REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	02/29/00	Final, 03/01/97-02/29/00

4. TITLE AND SUBTITLE
Computational Models of Human Organizational Dynamics
Final Report

5. FUNDING NUMBERS
AA,AB

6. AUTHOR(S)
Dr. Gregg Courand, Dr. Michael Fehling

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Synergia LLC
2400 Broadway, Suite 203
Redwood City, CA 94063

8. PERFORMING ORGANIZATION REPORT NUMBER
COD-13

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
DARPA/ISO
3701 N. Fairfax Drive
Arlington, VA 22203-1714

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release;
distribution is unlimited

12b. DISTRIBUTION CODE
A

13. ABSTRACT *(Maximum 200 words)*

14. SUBJECT TERMS
Organizational Dynamics,
Organizational Modeling, Organizational Analysis

15. NUMBER OF PAGES
61

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

Synergia LLC
2400 Broadway, Suite 203
Redwood City, CA 94063-1551
Phone: (650) 569-4999
Fax: (650) 569-4990
World Wide Web: www.synergia.com

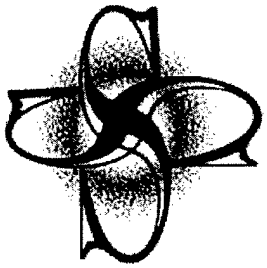

Computational Models of Human Organization Dynamics

Final Project Report

Sponsored by
Defense Advanced Research Projects Agency
Information Systems Office
Computational Models of Human Organization Dynamics
ARPA Order No. E495
Program Code No. 6S10
Issued by DARPA/CMO under Contract #MDA972-97-C-0001

Period Covered: 3/1/97 – 2/29/00

Executive Summary

This report documents the results accomplished by Synergia LLC for our research program on Computational Modeling of Human Organization Dynamics. This is the final report for our Phase II SBIR project, conducted over three years.

Our research program has contributed theory, methodology, and technology for organizational modeling and analysis. We have conducted successful tests of our work in a variety of domains, including crisis response planning, asymmetric warfare, municipal planning, health care delivery, and hazardous materials incident response.

We list, for each of theory, method, and technology for organizational analysis, the contributions made by our overall research program. We also describe organizational characterizations of important application domains. We describe each of these contributions in greater detail in the following sections.

Domain Characterization

Asymmetric Warfare

- This important domain is an extreme form of crisis. It deserves special attention because target populations generally have little understanding and capability to deal with the threat posed. Vulnerabilities derive from the interdependence of friendly weaknesses and aggressor beliefs about and capabilities with respect to those weaknesses; the particularity of this relationship mandates organizational analysis. Further, there are a very large number of targets and methods available to aggressors. Management of populations/assets under these uncertainties requires that we employ risk assessment methods to formalize decision-making to address the problem.

Crisis Characterization

- Asymmetric threat is an instance or example of a more general phenomenon, crisis. We have investigated crisis and crisis response in order to develop a wider view of the issues of asymmetric threat. We explored a variety of diverse crisis domains, and showed that our theory/ontology, methods, and technology provide very significant predictive power and support for the analysts and decision-makers in those domains.

Theory

Generative Practice Theory: Substantive Theory of Human Systems

- We have developed a theory of human systems, called Generative Practice Theory. It is based on the idea of "practices", which are adaptive generators of individual and

social behavior. With it we *describe* the physical and psychological characteristics of humans and the social structures of human organizations (ad hoc groups, teams, organizations, multi-organizations, cultures...). These descriptions give us a platform to *predict behavior* of subject agents. We have expressed it both computationally and algebraically, and we have developed a psycho-social ontology that embodies its commitments. We have used it to develop supporting theories of human personality and culture. It is the chief foundation for our methods and technology.

Methodology

Action Research: Overarching Methodology

- Our work with organizations is a type of Action Research. In other words, it is a science-based approach to the understanding and improvement of the capabilities of human organizations. Hypotheses are about the past, present, and predicted future behavior of subject organizations, under varying assumptions about environmental events and potential redesigns (interventions aimed at improving practices). Hhypothesis testing occurs through managed change experiments. This is an interventionist approach. Often it is a "dual intervention": our contributions are delivered as methods and technology to assist client organizations to intervene on themselves as a means to better intervening with others. For example, we can help a client see that the way it distributes decisions creates a vulnerability that an asymmetric threat could exploit – the adversary reveals something about itself. Then the client can use what it has learned to better manage itself and the threat.

Practice Mapping: Protocols for Acquisition, Validation of Data on Human Practices

- We have developed and tested a knowledge acquisition and modeling protocol, called Practice Mapping, to acquire the data on individual and organizational practices that is used to develop formal computational models of same. Highly ramified protocols for interviewing subjects have been developed. Instruments and observation methods have also been developed. We have also developed methods to gather data from secondary sources such as reports and the testimony of outside experts; this is necessary in domains (e.g., asymmetric threat) where the key individuals and organizations we seek to model are not directly available. All methods have been extensively and successfully tested, and all include validation internal to the methods.

Planning by Analysis: Central Problem-Solving Strategy

- We have invented a problem-solving strategy, called Planning by Analysis, that combines planning and decision-making with agent modeling and prediction/analysis for all participant agents, as well as information gathering to improve models. It is compatible with any existing planning technology (e.g., AI planners), and is optimized for complex planning situations in which agents face adaptive, intelligent adversaries. It includes methods and particular structures to support decisions about

how best to partition effort between planning based on what is presently understood, versus improving understanding as a basis for improved planning. It employs the concept of anomaly management and the methods of risk management within mechanisms that manage understanding, plans, and behavior.

Critical Dialogue: Multi-agent Coordination Discipline

- We have developed a coordination discipline, called Critical Dialogue. It is based on extensive review of crisis decision-making episodes (e.g., Cuban Missile Crisis), as well as leading theories in the Social Sciences. Critical Dialogue identifies and formalizes methods for individuals to disagree with one another in a way that is productive of superior joint decisions and plans. It posits a collection of dialogue roles; these roles are methods to manage the proposal, assessment, or repair of problem-solving elements (e.g., beliefs, decisions, models) and their inter-relations. Successful formal experiments were conducted at the Army's Command and General Staff College.

Technology

ACCORD: Technology Suites for Practice Description, Formalization, and Analysis

- We have prototyped a technology suite called ACCORD to support the description and analysis of individual and organizational behavior in terms of generative practice descriptions. ACCORD combines a prototype simulation driver[1] with technology suites for describing agent practices, converting those descriptions to computational form, designing simulation experiments, and analyzing data and results.

Consensus Facilitation Architecture: Distributed Technology Framework

- We have designed and begun to prototype an architecture, called the Consensus Facilitation Architecture, in which formal structures – e.g., decision models, ACCORD organizational models, or other models – can be proposed, shared, critiqued, revised, and employed by multiple users. It has a variety of special facilities to support collaborative real-time analysis, planning, decision-making, and other tasks. This architecture will also support wargaming studies – simulations that include computational and human participants.

In the remaining sections of this report we describe each of these application, theory, methodology, and technology development thrusts.

---

[1] The simulator has special features to support the execution of practices as we depict them (i.e., as action programs that are monitored concurrent with execution and redirected as part of execution to serve criteria). The prototype simulation driver was developed on another contract. It has been developed to execute over the internet and is being modified now to exploit a distributed network of processor resources.

Application Domain: Asymmetric Threat and its Management

In this section we characterize the asymmetric threat domain as a kind of organizational configuration. The terminology defined here is the foundation for an asymmetric threat ontology that integrates with our larger psycho-social ontology, presented later in this report.

We also present an analysis method for managing asymmetric threats. It provides for:

- discovering weaknesses in friendly organizations and facilities,
- developing intelligence collection processes to detect aggressors pursuing asymmetric methods,
- redesigning friendly organizations,
- evaluating friendly plans to thwart aggressors (pre-emptive and responsive).

We sum up with requirements for support of the management of these threats, and present a few brief remarks on the present state of the art with respect to these support requirements. This is followed with a collection of more general remarks on crisis, of which this domain is an example.

Asymmetric threat refers to threat configurations in which one agent, an **aggressor**, employs weapons against a relatively defenseless asset. **Assets** are associated with one or more friendly agents, and have a variety of forms: populations, infrastructures, knowledge, capabilities, symbolic value, etc. Asymmetric threat is contrasted with the classical conception of warfare as force-on-force engagement between roughly equal agents. Terrorism (physical, biological, and cyber), and some types of insurrection are examples of asymmetric threats.

To develop an approach to managing asymmetric threats, we distinguish between "weaknesses" and "vulnerabilities". A **weakness** is some condition of an asset that could conceivably be exploited to harm or destroy the asset. In this sense, almost every asset has a large number of weaknesses. An asset is **vulnerable** only if it has a weakness for which there exists an aggressor capable of a) detecting/recognizing the weakness, and b) exploiting it to serve aims of the aggressor. A **target** is an asset or a feature of an asset that an aggressor is actually considering attacking.

These distinctions allow that there are weaknesses that are not vulnerabilities, since there may in principle but not in practice be an agent capable of detecting and exploiting the weakness. They also allow that a target may or may not correspond to a weakness, and even it if is a weakness, it will only be a vulnerability of the asset if the attack capabilities of the aggressor are sufficient.

One consequence of this view is that, unlike classical warfare, with its long-term development of methods and weapons, asymmetric warfare is usually highly particular. *Asymmetric warfare depends on the specific understanding, aims, and capabilities of the aggressor* in relation to *the weaknesses actually exhibited by an asset.* For example, the assassination attempt on Egypt's President Mubarak by Islamist extremists exposed particular weaknesses in his security

protocols, as well as particular strengths of the operatives. (Luck more than anything else saved his life.)

Now let's consider the assassination attempt on President Mubarak, but from the point of view of the role that could be played by our methods and technology. This case can be modeled carefully, and then the practices of his security forces can be redesigned a) to remove the weaknesses exposed by the attack, and b) to implement a COA to interdict or destroy the threat. Further, this type of attack can be generalized by varying the resources and methods of the aggressor – until we finally have a family of aggressor methods for attacking the motorcade. Then any similar target (e.g., U.S. Presidential motorcade) can be evaluated for weaknesses by presenting it with the challenge offered by the family of aggressor models. *In general, within an analysis of vulnerabilities we can predict aggressor behavior, discover friendly weaknesses, and evaluate candidate repairs.*

This example, and the generalization to families of aggressor models, is meant to suggest a method for managing asymmetric threats. The method presented below is an instance of Planning by Analysis, because it centers on the role of modeling and then predicting/analyzing the result of simulated asset/aggressor interactions. Analysis continues with the proposal of candidate redesigns (of the assets physical and organizational features), as well as their COAs, and concludes when results are satisfactory.

We present the asymmetric threat management method here in its central logical form. But as we go on to describe below, it is best implemented within a risk management discipline.

1) Collect a library of asymmetric warfare incidents (e.g., assassination attempts, facility bombings, hijackings, biological and chemical weapons use, etc.)
    a) Classify the assets, agents, and targets involved.
    b) Agents include aggressors, organizations associated with assets, as well as friendly organizations charged with responding to, interdicting, pre-empting, or otherwise thwarting aggressor actions against the assets.
    c) Develop agent models for assets and aggressors – practices and their coordination within particular organizational structures
    d) Validate agent models (e.g., postdictive and predictive studies).
2) Develop cases for prediction and analysis.
    a) A **challenge** is a postulated aggressor agent, endowed with given resources, and endowed with given methods that are coordinated within a particular organizational structure
    b) Produce a set of types of challenges by generalizing over known experience with aggressor organizations
    c) Produce a set of assets that will be treated as targets. An **asset** is the combination of a physical infrastructure and an agent model that describes the asset's associated practices and organization.
        - Assets might be ones of special interest. Alternatively, an asset set might be formed by generalizing over the actual targets chosen by aggressors, until the generalizations cover friendly assets of interest

d) When relevant, include as participant agents any organization that will execute a COA to thwart the aggressor.

e) **Cases** are combinations of challenges, asset models for friendly targets, and agents that execute COAs against aggressor agents

3) Simulate and/or wargame the cases to predict the outcomes – e.g., for the friendly asset, the agents executing the COA, and the aggressor.
   a) Determine what counts as acceptable performance for the asset and its plans.
   b) Determine the independent variables – states of environment, target, and challenge that are to be systematically varied.
   c) Determine the dependent variables – the measurements to make. Usually, these will rule on hypotheses about asset performance or COA effectiveness. In most cases, there should be attached validation tests.
   d) Collect the results for analysis – usually, in the form of probability distributions over outcome variables.
4) Analyze the results to diagnose weaknesses in friendly assets and plans
   a) Determine the causes of adverse outcomes of sufficient probability by tracing outcomes back through prior event trajectories
   b) Localize primary causes to practices by determining the sensitivity of final, adverse outcomes to particular prior event trajectories
   c) Summarize weaknesses: annotate event trajectories with the weakness in the asset (organization and plan) and/or COA that is opposed by the strength in the aggressor (organization and plan).
5) Define and redress information shortfalls associated with weaknesses.
   a) Test to see if important information is associated with weaknesses – i.e., a weakness has been discovered, and part of the asset model or COA or aggressor model contains highly uncertain elements.
   b) If uncertainty is significant, and if possible, then gather relevant information and refine agent models.
   c) Rerun the cases that exposed the weakness under the refined agent models, to see if any weaknesses have been eliminated (new ones may also appear).
6) Determine the degree to which the modeled threat's characteristics are embodied in actual agents.[2]
   a) Develop indicators for the practices of the modeled aggressors.
   b) Review intelligence data and implement intelligence collection programs to determine whether observed behavior of individuals and organizations is consistent with the modeled aggressor's structure and method (including its preparation).
   c) Based on the intelligence results, evaluate the probability that there exists a real aggressor that could behave as the modeled one (that exposes a weakness).
   d) A **vulnerability** is a weakness for which the probability that an actual aggressor either exists or could reasonably be expected to come into being is above some threshold.

---

[2] Recall that the aggressor models are generalizations from experience with actual events. As such it is not certain that an actual aggressor exists that can successfully harm the friendly asset. In all likelihood, the key is to analyze whether an aggressor of this form could come into being.

7) For each vulnerability, propose and model practice redesigns (intervene on self) as well as preemption and response plans (intervene on aggressor)
    a) Develop friendly courses of action to thwart the detected real aggressors.
        - Also define any requirements the COA places on the asset (e.g., changes so that the COA can be executed.
    b) Develop practice designs adequate to implement the COA and/or repair the asset to remove the vulnerabilities (weaknesses that the aggressor exploits). Such designs may be any combination of:
        - policies and practices
        - organizational structure
        - physical infrastructure arrangements
    c) Gather relevant information to refine asset and agent models and/or the COA.
8) Return to step 3, but this time, evaluate the cases under the proposed redesigns. Iterate through the next steps until satisfied.
9) Implement effective redesigns.

It is clear from this proposed method of analysis that significant combinatorics may need to be managed. For any given asset, it is necessary to consider the potential threat posed by numerous potentially-relevant aggressors (or generic types of aggressors), perhaps using a variety of methods (e.g., car bomb, ultra-light delivery of chemical weapon, etc.).

Also, a risk management discipline is needed to make use of the results of this analysis. At one level, the asset/aggressor combinations are not all created equal. Some weaknesses are more likely to be vulnerabilities than others (i.e., real aggressors really interested in the asset), and some assets are more valuable than others. Risk management (i.e., probability and value modeling) is a proven framework for quantifying these considerations. At a deeper level, we believe it is likely that there are many more vulnerabilities than there are resources to repair the underlying weaknesses. If this is true, then risk assessment is essential since it provides a normative basis for deciding which weaknesses to repair, in what order, and to what degree. In this way risk management provides principles and tools for determining which cases of the above analysis method should be looked at, what information should be collected, and what responses are warranted.

In sum, we believe the following elements of method and technology are required to support organizations charged with managing (understanding and dealing with) asymmetric threats.

1. **Psycho-Social Ontology and Modeling Capability.** It is necessary to be able to model the views, operating procedures, capabilities, and aims of agents (both threats and aggressors). In this way the complex of weaknesses and perceived-weaknesses (i.e., vulnerabilities) can be properly described and analyzed. This is essentially a requirement for a modeling language or ontology that captures psychological and social/organizational structure, and for a suitable method of analysis.
2. **Complexity Management Technology.** It is necessary to evaluate the interaction of assets and aggressors under a wide range of conditions – ranging from variation in stochastic input events and randomized outcomes to variation in the high-level perspective, aims, and abilities

of agents. This is essentially a requirement for a comprehensive simulation (or wargaming) and analysis capability.

3. **Risk Assessment.** Asset damage occurs when an aggressor recognizes and then exploits a weakness of a given asset with a given method. The actual amount of damage is uncertain and best depicted as a conditional probability distribution, i.e., damage is a random variable that depends on the asset, and the aggressor's choice processes (targets and attack methods) and capabilities. The valuation of asset damage derives from some quantitative judgement about the loss of worth of the asset as a function of the type/amount of damage. These two kinds of quantification – probability of damage and loss of worth – are properly integrated by formal risk assessment models. Then each of the steps in the above management method are either information-producing experiments or are decisions about what is best to do next.

We have reviewed and previously reported on the state of the art relative to these requirements. The state of the art in psycho-social description/ontology has been relatively primitive – as compared with exclusively physical, psychological, or social characterizations. However, we have developed a spanning psycho-social ontology (GPT). Similarly, there is relatively limited simulation or wargaming support for psycho-social processes (e.g., simulators like MODSAF depict physical objects like tanks, but not the minds in the tanks, nor their coordination). Our own work on ACCORD redresses this deficit and encapsulates analysis within experimentation. There are a variety of promising risk assessment technologies in the form of decision modeling and analysis environments. In other research we have developed a method to integrate psycho-social modeling, extensive simulation, and formal (decision-theoretic) risk analysis. This method offers a rational (decision-theoretically sound) basis for determining vulnerabilities and assessing plans for repairing weaknesses and interdicting aggressors.

Widening our Focus from Asymmetric Threats to Consider Crises in General

The domain of asymmetric threats and their management has been the chief area of concern for this project. However, it is important to test our concepts, methods, and technologies on other, similar domains, since they will test our framework in different ways, reveal different opportunities for its use, and etc. In this section we present the domains we investigated, our results in those domains, and then a general characterization of crisis.

We have looked at a wide variety of domains that in general can be grouped under the heading of "crisis" and "crisis response". The primary domains we investigated, and the primary result of each, follow.

- South Africa had a covert program to develop nuclear weapons, and then later a mostly-covert program to destroy the weapons and dissolve the original program. We had studied it during Phase I of this SBIR; we reviewed and improved our analysis of the case during Phase II, using our more powerful theory and ontology for organization. *We showed the significant predictive power of a combined key-figure and institutional analysis. Experts on the case were compelled by the analysis.*
- The Environmental Protection Agency responds to and manages hazardous materials releases. They also have an ongoing activity to deploy response resources against need, and helping communities develop response organizations. We carefully

modeled the planning as well as the response organizational dynamics. *Our model of this process was judged superior to the one the inventors had spent two years developing. They decided to consider it as the basis for their training program.*

- Health care delivery manages crises in an ongoing way in Emergency and Critical Care Units. We modeled the delivery of services in the Critical Care Unit at Stanford Hospital. *We elucidated important features of doctor-patient interactions that should figure in redesigns of the CCU.*
- We developed a crisis response planning method (Critical Dialogue), based on our theory of social/organizational work, and an extensive review of transcripts of dialogue during important crises (e.g., Cuban Missile Crisis). We trained a collection of Army officers at the Command and General Staff College in the methods. *In formal experiments the trainees produced significantly improved crisis response plans to hypothetical problems, as compared with the untrained control group. Their plans were superior, and review of their reasoning showed they tended to generate more options and improve their analysis of the most compelling options.*
- We developed a high-fidelity scenario of an insurrection and located it in Dar Es Salaam, Tanzania (prior to the embassy bombing). We used our prototype technology to illustrate real-time agent modeling, planning, plan monitoring, information gathering, and replanning to conduct a non-combatant evacuation operation. *We demonstrated the potential of a mature technology of this type to mitigate or eliminate the biggest concern of such operations – as explained to us by experts at SOCOM. Specifically, we can detect when troops either have begun with very faulty information, or when the behavior of the adversaries has changed in a way that is significant for the current COA. This leads directly to improved real-time, in situ planning/replanning.*

Our Characterization of Crisis and Crisis Response

We have conducted an extensive review of the theoretical literature on crisis, as well as the literature on specific crises (e.g., Cuban Missile Crisis, Mayaguez Incident). In contrast to most of what we find in the literature, we have found it useful to distinguish crisis, crisis response, and crisis response planning.

Agents (individuals and organizations) experience crisis whenever the combination of threat, urgency, and uncertainty are such that the agent feels compelled to hazard critical change. In other words, crisis is revealed by significant and important changes that compel us to act, even though the limitations of our understanding mean that this may significantly worsen the situation or our future response capacity.

We treat **crisis** as a description of human cognitive state, and have distilled its dimensions to three: threat, urgency, and uncertainty. This characterization can be understood as follows. Agents may classify their situation according to: the risks it poses to goals or more broadly to values, the capacity they have to achieve their goals, and the quality of their understanding of the situation. When risks exceed some threshold of tolerance, they become threats. When an agent believes a response is necessary (i.e., "must do something"), but the agent's capacity to respond,

given available resources and agent skills, falls short of what is needed in the current context, then the agent experiences urgency. When an agent's understanding falls short of what is needed to adequately characterize the situation or (especially) to create and undertake a suitable response, then the agent experiences uncertainty. A crisis occurs as a function of these three elements. That is, an agent is (defined to be) in crisis whenever the threat, urgency, and uncertainty, either individually or in some combination, exceed some tolerance threshold – so that the agent feels compelled to hazard critical change.

Our definition of crisis embodies the following commitments. First, the proper measure of urgency is agent capacity as it may be deployed in time, by specific agents, in specific circumstances. Urgency is not defined by available time per se, and, what is urgent for one agent might not be urgent for other agents, and might not be urgent in other circumstances. An event projected to occur in 100 years, but for which we have no adequate response (e.g., catastrophic global warming) is sufficient for crisis, whereas nothing counts as a crisis, no matter how bad, if we have or will have the capacity to handle it. Second, uncertainty applies to the detailed parameters of a situation, to the distinctions that are relevant to apply (explanatory categories), and to the decisions and actions that are appropriate. Third, each of these dimensions of crisis can be developed normatively, that is, with respect to a rational decision-making formulation.

Crisis response is any pattern of activity, some portion of which is intended to positively affect any of the dimensions of crisis. Crisis response support is that collection of methods and technology that are developed and deployed to predict crisis and/or improve our ability to manage our responses to its potential or actual occurrence. Based on our experiments to date, we believe that psycho-social modeling within a risk management discipline, and employing a supporting technology, will prove effective for crisis response for managing asymmetric threats and for each of the crisis domains we have investigated.

Generative Practice Theory

Generative Practice Theory (GPT) is our substantive theory of social dynamics. It is based on an extensive review and synthesis of leading thinking in Psychology and the Social Sciences, and forms the foundation for all of the social/organizational claims and supporting technologies developed during this project.

Generative Practice Theory's main commitments are to the concept of practice as the generator of adaptive behavior, to an agent architecture, and to a view of social structure as arising from regularities in behaviors generated by co-adapting practices. In GPT, social systems are not just represented computationally – they are a kind of computation.

We discuss GPT in detail in the sections that follow. We conclude the discussion by summarizing GPT's commitments as a **psycho-social ontology**. This ontology forms the basis for rigorous, analytic modeling of individual and social capacities.

Practice: the Base Construct

We explicate our view of practice as follows. When we observe humans and their organizations, we observe behavior. We observe that similar behaviors are enacted, with similar results, in similar situations. In our development, practices are the generative mechanism responsible for producing behavioral outcomes in the context of given situations.

Agents rarely if ever respond to a situation that exactly matches a previous one, yet their behaviors achieve similar results over a range of circumstances. As an important and revealing special case, agents correct errors to bring performance back to desired outcomes. This reveals that agents are fundamentally adaptive. We claim that the regularity in their behavior and results are the result of striving to satisfy stable (for some time period) collections of criteria. This is what it means to be skillful, or have a competence – behavior attends to and seeks to satisfy criteria.

Behavior creates outcomes, such as changes in physical conditions and/or cognitive state. Additionally, behaviors are observable to the agent and other agents to a greater or lesser degree, in this way contributing informational ingredients for other behaviors. Agents have access to their behavioral products, and those of others. Stable relations in behaviors, describable as being generated by joint practices, are the ingredients for our theory of social structure.

In this theory, human beliefs, desires, intentions, feelings, and etc., both individually and socially, are embodied in capacities to behave. For example, to orient behavior to exploit and damage a given feature of some asset entails the belief by the agent that the feature is a weakness of the asset. Note that it is not necessary to have articulated the belief; it is only necessary that the action orientation is consistent with the belief. It is not hard to see that almost all human belief is tacitly embodied in practice; only a small part of it articulated.

In general, the situations in which it is relevant to produce a behavior, the criteria that guide adaptation, and even the general program of action, all tend to be tacitly embodied, not articulated. Description/articulation, when present, is the product of reflective capacities.[3] This accords with psychological evidence.

With these considerations in mind we define a **practice** to be the generative capacity to produce a behavior, in a class of situations, and subject to a set of criteria (i.e., conditions to achieve, preserve, or prevent). **Behavior** is the manifestation of enacted practices: practices when enacted generate adaptive behavior and thence individual and social products.

---

[3] Our approach contrasts with the view that humans think in propositions. Gilbert Ryle introduced the distinction of "knowing how" versus "knowing that". A substantial amount of AI and other research treats humans as 'knowers that' (i.e., as having minds composed of propositions that get manipulated). We believe this is flawed, even if one wishes to build a classical expert system. But in any case, it fails as a description of real human individuals and societies. To know is, for this theory, to know how – to be able to generate a behavior in a regular way, with regular results, under a range of execution conditions. Propositions are abstractions on behavior.

Practices are a formal construct that we are introducing to formalize descriptions of adaptive, situated behavior. Since we are developing computational models of behavior, we sometimes find it necessary to distinguish practices, the formal generative object (e.g., represented as computer programs), from dispositions, the actual behavior-producing capacities resident in human beings. We simply point out the distinction here and leave it as understood for the remainder of this report.

It is a central strength of this view that a practice is sufficient to encode primitive skills, complex skills, individuals, organizations, or any other social system. We will see this in more detail below; however, the key point is that agents and social structures such as organizations can always be represented as clusters of practices that exhibit particular and regularly-repeating coupling relations.[4]

We represent practices in three languages: computationally, propositionally, and algebraically. The propositional and algebraic representations are derivative of the computational one.

The computational form of a practice, called an **action schema**, is defined by four components:

- **Procedure:** an action plan; it generates the behavioral products of the practice. The specific products depend on information available at the start and during execution. Elements of the action plan are implemented as continuations, allowing for the possible restarting or redirecting of the procedure under revised conditions.
- **Trigger:** conditions, in the form of a pattern of events, that are necessary for the procedure to be a candidate for execution. They may not turn out to be sufficient to fully warrant execution; higher-level practices may manage triggered ones according to the situation, available execution resources, and etc.
- **Binders:** information on which the procedure and monitors depend or are influenced. May be produced by the procedure. May also access part of the external context of the action schema, agent, or external environment (i.e., make external events visible). Frequently, there are antecedent events that are not direct triggers, but that nevertheless shape the way the practice executes.
- **Monitors:** active tests of the criteria associated with a practice. They monitor patterns in the binders; they operate concurrent with procedure execution. When they detect a violation of the criteria associated with the practice, they suspend, abort, or redirect the procedure.

An agent's attention is defined by the set of action schemata whose triggers and monitors and procedures currently have execution resources. This determines an important relationship between an agent's action schemata and the overall architecture of the agent.

---

[4] We have claimed that human knowledge is best understood as practical, not propositional – knowing how, not knowing that. Propositional accounts necessarily need to treat social belief (and value, and etc.) as relations over individual belief, and thus as constructs of a fundamentally different order. This leads to the mistaken view that 'the social' is merely the composition of multiples of 'the individual'. Social practice in GPT is co-adapting constituent practices, and so is just another practice.

The propositional form of a practice, called an **action proposition**, is the declarative form of the action schema. It is described by the formula In Situation S, Do Action A to Satisfy Criteria C. This form is used by analysts wishing to describe practices (e.g., users of the ACCORD mapping technology). The situation S describes all situations that may influence the execution of the action plan, and so includes the triggering situation, others that may arise during the execution of the practice, and final products or results. The criteria C describe the conditions the practice seeks to produce, preserve, and prevent. The action A is just a statement of the (advisable) action plan – the plan that is specialized by the criteria according to the situation.

Practices have been represented algebraically as **Bucci automata,** and more recently, using a variant of the **π-calculus** (Milner). In each case, primary concern is for a formal expression of interruptible, non-terminating computation. A Bucci automaton may read an infinite sequence of characters on its 'tape', and is defined by its response to patterns in those characters, rather than to each character as it is read. The **π-calculus** additionally offers a more powerful capability to formalize concurrent practice.

Agent Architecture: Practice Management via Ambiguity Reduction

Practices may contain practices – action schemata may contain action schemata within their action plans. From this it is possible to build essentially any type of control regime for assembling and managing the activity of constituent practices. Ultimately, this means that an **agent**, be it an individual or some social entity, is composed of structured bundles of practices. Even so, it is necessary to have an organizing principle that determines how practices should be structured and, especially, how their joint execution should be managed. Simply put, an agent needs an architecture to answer the question of 'what to do next'.

We have developed a preliminary specification of an agent architecture, called the **Ambiguity Reduction Architecture**. Ambiguity arises whenever agents have a task but lack an executable action schema. They may lack a suitable practice, have multiple candidates rather than a single one and so face a choice, or have conflicting practices. The principle used to structure and manage cognition is, roughly: reactively execute practices whenever they are fully executable and execution resources exist; otherwise, deliberate to resolve whatever ambiguities are blocking reactive execution.

We identify deliberation (as opposed to reaction) with the practices that resolve ambiguities and so create the executable action schema. We distinguish two main types of deliberation: selecting among candidate practices, and constructing a suitable practice. Selection practices resolve conflict; they implement a choice process over multiple candidates (alternatives that may or may not conflict) and produce a partial order over a non-empty subset of the candidates.

Constructive practices produce a suitable practice from available ingredients. Sample constructive procedures include:

- instantiating variables in otherwise-fully-specified practices, according to elements in the current situation, thereby creating the executable practice

- assembling a practice from elements of other practices, e.g., this action plan but serving these other criteria
- mimetic operations that attempt to produce a practice that generates a copy of the behavior of another agent.

Selective and constructive strategies may be combined, e.g., select among alternative policies for action, and then instantiate the selected one with information about the current context so that it is ready to execute.

We are undertaking preliminary steps to assess the empirical validity of the architecture. In particular, we are devising data collection protocols that investigate the amount of preparation (i.e., deliberation in support of selection and/or construction) that individuals and groups engage in to produce an executable practice.

Social Structure in Generative Practice Theory

Broad Features of Social Structure

We identify **social structure** in Generative Practice Theory with patterns of joint/multi-agent practice that generate recurring dynamics of behavior. Structure, for this theory, is structure-of-process. GPT makes the claim that there are exactly three dimensions of social structure, namely, **resource relations, communication relations,** and **role relations**. We describe each below. We begin general remarks on social structure – its nature, and how it arises.

Social structure emerges and is perpetuated because agents assimilate and learn and refine practices by adapting to what they themselves are doing, what other participants are doing, and (especially) the combined/joint behavior that all participants are producing.

Concretely, agents adapt their behavior as participants, according to their view of (their relation to) the overall system of behavior. Consider, for example, the social act of two people carrying a table through a door that is just big enough to allow the act – many readers will have experienced this or something similar. In this case, each person is guided predominantly by the state of the overall system – mainly the table and where they are in the overall process – at each moment. We note that they attend only to their own or the other's detailed state when a problem arises – in fact, this is the evidence that suggests the participants tend to operate with respect to the social state as a primary orientation. Similarly, consider the social act we call "a meeting". Every such act is unique, yet participants tend to manage their individual behavior to so that when combined with the behaviors of others the overall product conforms to something the participants recognize as "a meeting."

How can this work? Are we claiming that everyone has the same model of a meeting? Yes and no. Yes in the sense that there is a family resemblance in the pattern of actions, and a collection of 'rules' that determine how individual contributions are introduced. No in the sense that we do not think the image of the overall pattern is identical. It does not have to be; agents are adaptive!

How then do the similar understandings arise? They arise as the learned experience of being in the joint practice. Any time we enter a social act, we are both employing and perpetuating our understanding of this social act – making it possible for important social practices to emerge, stabilize, and be learned by new agents. And so the reality of social structure lies in the ways we jointly construct and perpetuate joint actions. Moreover, it does so without anyone needing a definition or explicit understanding of the act. It is enough that agents act in ways that jointly orient toward the pattern. This view of social structure is a form of cognitive institutionalism.

In terms of our definition of practice, social structure is determined by stable configurations of criteria in the constituent practices – the social 'control law' is produced by these stable configurations. As may be seen in the table and meeting examples, we hypothesize that the criteria sometimes refer to properties of the overall system. This allows agents to adapt their behavior in the light of the total state, and so help move the total state in a coordinated way through their individual actions. If the criteria reflect similar overall states then the individual acts will combine harmoniously and the total system behavior will conform to the several criteria.

And so we are led back to the question, how then do criteria on the social practice emerge? In many cases, social structure is perpetuated, within a population or culture, through a combination of mimetic and adaptive behavior. Consider the new participant (referred to by Lave as peripheral). This participant will first experience the system as a collection behaviors that attend to criteria, and then will learn over time the ways of 'fitting in' to the system. And so individual adaptation begins with the rules that shape the social act.

It may be noted that this view of social structure is entirely 'internal' to systems of co-adapting practices; in part, this is just what it means to be a generative theory. Yet a social scientist might complain that a lot of what might be seen as social structure is summarized by what is not done – either due to prohibition, or more commonly, due to the behavior not being conceivable. Families, tribes, organizations, and cultures have options that are in principle available to them, but in practice are never realized because they are not imagined. For example, the injunction against members of some native communities from leaving their tribe, which operates as if it is a law, is actually embodied by having no practice for conceiving and manifesting this 'option'. This type of social structure is revealed by attempting to change the system (perhaps only as a thought experiment) and investigating how criteria-driven behaviors guide the system to preserve its form – i.e., by studying its adaptive response as a function of actual practices. And so once again, social structure – including its frontier of never-enacted options – derives from (is generated by) a system of practices.

Dimensions of Social Structure: Resource, Communication, and Role Relations

By resource relations we mean regular patterns of adaptive behavior that can be analyzed into stable criteria for accessing and sharing physical and informational constituents employed by a practice. Resource flows are revealed by the physical and informational constituents that are typically employed when a practice executes, and especially by the criteria that determine who, when, how, and how long scarce resources are accessed, consumed, created, transferred, or otherwise shared among practices and agents. Resource structure is also revealed by the

'standard interpretation' given to physical spaces and artifacts (e.g., this space is a conference room, and thus endowed or invested with uses deemed standard).

Communication is not equated in our work with information (data) transfer. We treat information transfer or access as a type of resource access and utilization. Instead, communication is fundamentally about the relationship between what is said by one agent and what is said in response by that agent and by other agents. Thus, communication is fundamentally about how interpretations are formed and how they guide later action, especially, speech.[5]

Communication centers on the practices of interpretation that give meaning to whatever information happens to have come to the attention of an agent. "Meaning" here refers to changes in the disposition agents have toward one or another practice – meaning is the change in the commitment to action entailed by processing information. Concretely, it is any change in the conditions that trigger or that might direct or redirect the execution of a practice. Given this, by communication relations we mean any regular pattern of formation or dissemination of interpretations, via communications actions (linguistic and non-linguistic). To simplify, it is the relation between received input signals (information as resource) and output responses in the form of speech production and/or other acts.

We organize our analysis of communication relations in terms of topics. Topics are templates that tend to be called forth in given communicative interactions (e.g., conversations) and that inform the information environments of active agent practices. Topics determine what information will be selected (from whatever is available), what sense the agent will make of what has been selected, and the implications for later actions (especially, communicative response). Topics are very useful as summaries of the varying perspectives agents tend to bring to any joint action.

Roles and role relations are explained as follows. Practices are not independent/autonomous bits of skill; rather, practices are bundled together (directly, and in terms of the situations to which they respond) as practice complexes that maintain a collection of interdependent criteria. These inter-practice complexes are roles. For example, what we think of as a profession or as a recognizable social entity (e.g., doctor, parent) is not a flat collection of practices. Instead, it is a complex of practices that are coordinated to achieve a range of criteria, standing in particular relations (e.g., standard trade-offs), enacted in a given range of situations and tending to embody a given set of methods.

Roles, strictly speaking, are resident in the social system of interest; individuals 'play' their roles, thereby giving it meaning, receiving censure or reinforcement, and potentially changing what is meant by the role in the overall social system. Thus roles emerge and achieve stability by co-evolving in relation to other roles. To be a doctor is to evolve a stable pattern of skills by being confronted with doctor~patient, doctor~doctor, doctor~medical-system, and other doctor~agent encounters. To be a parent is to evolve a stable pattern of behavior through repeated

---

[5] We do not treat this in detail, but non-linguistic actions such as gestures have interpretable symbolic content, and so we can treat them as communications, just as we do for speech.

interaction with children, spouses, one's own parents, other parents, and so on. In these examples, doctors and parents are learning their roles through their ongoing experiences of the joint encounters with other agents playing complementary roles. Hence to count as a role we must find a role relation, each role of which can be assumed by multiple agents.

Roles are explicated in terms of the expectation participants and stakeholders have for the tasks and performances associated with the role. Tasks are requirements and constraint specifications on admissible and desirable procedures, agents who may or may not participate, and expected results. By analyzing the expectations agents have for the tasks an agent sets and/or responds to (for themselves and for other agents) we determine the roles. By analyzing the way tasks are partitioned among or shared over agents, we determine the role relations.

Role relations establish the large-scale governance of practices across agents; resource and communication relations are constrained by the content of role relations. Social constructs such as authority relations, power or influence, and legitimation can be analyzed in terms of role relations, with particular attention to the ways agents segregate and control resources and communicative flow. One way to do this is to extract the decision-making purview embodied in a role relation, and then derive the implications for the flow of resources and the communicative tendencies among agents. Authority (and power, and legitimation) are thus analyzed by the content of decisions each agent is empowered to make, the practices that sustain decision-making purview, the way decisions are localized to or shared across agents. This pattern of decision-making authority in turn shapes interpretation, conversation initiation, and then the information that tends to be desired and that tends to be available when making each decision.

Psycho-Social Ontology: GPT as a Formal Language

This section summarizes the foregoing discussion of GPT as an ontology, sufficient to express the content of any psychological or social system. In addition, it has the main categories to express physical objects and processes. In this way, the extensive work on ontology in the Artificial Intelligence community, which focuses almost exclusively on the representation of physical categories (including information) can be integrated into this spanning ontology.

We present the ontology, proceeding in the standard way from foundational categories to the full specification of social/organizational systems.

Foundational Categories

- **Thing**
  - — All objects, properties, relations, etc. are Things
  - — Cf., KL-ONE and Conceptual Graph Theory (Sowa)
- **Relation** *ako* Thing

- association among Things

- **Property** *ako* Thing
  - *has* Value(s) (= Thing(s))

Primary Physical Categories

- **Duration** *ako* Thing
- **Time** *a set of* Relations among Durations (n.b., the universal resource)
- **Event** *ako* Thing
  - *has* associated Disposition
  - *has* Duration (I.e., location in Time, temporal extent)
  - *consumes, produces, or modifies (properties of)* Things
- **Object** *ako* Thing
  - *has* Structure (relation(s) among constituent Things)
  - *has* Physical Extent (I.e., (distributed) location)
  - *has* Duration

Primitives for Behaving Objects

- **Resource** *ako* Thing
  - *has* associated Disposition(s)
- **Disposition** *ako* Thing
  - *has* Locus (= associated Object(s)
  - *has* Initiator(s) (= Events sufficient for "activation")
  - *has* Monitor(s) (= embedded Event- (interrupt-) responses)
  - *has* Procedure(s) (= spec of flow of events)
  - *has* Resource Constraints (= *a set of* Resources necessary for "execution")

- *has* associated Effects (= Resources produced or modified by Disposition)

Actors: Control Regimes for Devices and Human Systems

- **Actor (= Dynamic Object)** *ako* Object
  - *has* Receptors (sub-objects associated with Initiators & Monitors)
  - *has* associated Disposition(s)
  - *has* Control Structure (I.e., decision policy for conflict among activated dispositions)
  - *has* Effectors (objects associated with potential effects)
- **Action** *ako* Event
  - *has* associated Actor(s), Disposition(s)

Agents: Actors with Purpose (strictly Human Systems)

- **Purposive Disposition (= Practice)** ako Disposition
  - Criteria (events/conditions) preserved/managed by Monitors or Procedures
  - *Intent*: associated criteria
- **Agent** *ako* Actor
  - *has* Purposive Disposition(s)
- **Intelligent Agent** *ako* Agent
  - *has* Practices to modify, construct, or remove Practices

Primitives for Social Structure

- **Resource Relation** *ako* Relation among *Producers* and *Consumers*
  - *has* associated Resource(s)
  - *Producers* (= actors whose effects produce the resource)

- Consumers (= actors whose practices are constrained by the resource)

- **Role** *ako* Thing
  - *a relation among a set of* Practices

- **Role Relation** *ako* Relation
  - *has* associated Roles

- **Communication Skill** *ako* Disposition
  - *has* Signals (= associated events)
  - *has* Interpretations (= Practices whose Initiators are matched by Signals; may produce Signals)

- **Communication Relation** *ako* Relation
  - *has* associated Topic(s)
  - *has* associated Resource(s)
  - *has* associated Agents who send, receive, and interpret signals

Social Structure

- **Organization** *ako* Actor
  - *has* Participants (= associated Actors and Agents)
  - *has* Resource Relations associated with Participants
  - *has* Role Relations associated with Participants
  - *has* Communication Relations associated with Participants

- **Organizational System** *ako* Organization
  - *has* Organization(s) as Participant
  - *has* Environment (= associated objects, events, and effects)

Practice Mapping

Introduction

In essence, Practice Mapping is a knowledge acquisition and empirical studies discipline, used to gather and validate data about organizational practices that is in turn suitable for modeling and analysis activities. **Practice Mapping (PM)** is a collection of protocols we have invented to develop and validate practice, agent, and organization models as specified by GPT.

We refer to the individuals that collect the data as **mappers**, and the collected data as a **map**. We refer to the individuals that offer their views of behaviors as **client-subjects**; clients because they rely on mapper skill and leadership of the modeling exercise, and subjects because their testimony is treated as a collection of hypotheses specific to each respondent and to be tested by mappers.

We have developed a very extensive collection of methods, instruments, and associated tools and technologies to facilitate mapping. We have conducted extensive research and testing of our methods, and have drafted a book-length monograph on them.

In the sections to follow we describe highlights of this collection of methods. Our presentation is meant to be general across any domain of organizational activity, and so the methods are explained in abstract terms. *We add that at the time of writing of this report, we are (on another project) testing a concrete step-by-step mapping protocol. It is based on the material presented in this report, but can be carried out by a mapper with minimal training and no special background in organizational analysis or human psychology.*

Sources

When working directly with an organization, we employ interview methods with members of the organization and knowledgeable outsiders. We also observe practices (e.g., as do ethnographers) within physical environments – i.e., in relation to the physical plant. And, we have been very successful developing and employing specially-designed instruments that determine the key practices to focus on, and that develop the organizational structure relevant to those key practices.

In some domains, such as the asymmetric threat domain, direct interviews and observations with members of the subject organization is difficult or impossible. In these cases secondary material and experts are our primary data sources. For example, we have mapped substantial elements of behavior of a particular asymmetric threat agent of global interest. We have studied open-source literature and other documents, and we have interviewed experts that have a detailed experience with the practices of the subject agent – in our case, FBI profilers and intelligence officers. We have established to our satisfaction that is a viable method for producing valid, predictive models of the subject agents.

Facilitating the Articulation of Experience

Agents (individuals, groups, organizations) rarely have more than a minimal pre-existing articulation of their practices. Usually, they can produce narratives that describe compelling features of their experience of their own and the behavior of others. As such, PM is a process of facilitating the articulation of experienced and projected behavior (and in particular, its articulation in the light of the GPT distinctions used to develop the associated practice models). It includes careful testing of the described behaviors and of the mapper's understanding of the descriptions. It is not an interview technique that merely asks for information that agents/respondents already have on hand.

Mappers need a collection of special dialogue skills, mainly because practice mapping is a process of facilitating articulation. In addition, social phenomena can be very subtle, subject to widely varying interpretation, and protected by defenses. We do not discuss these skills here, but, mappers need to be skillful at productively: managing conflict, uncovering and probing biased description, testing their understanding without 'putting words in the mouths of client/subjects', recognizing and testing their own assumptions, testing whether they have talked to the right client-subjects, recognizing when undue weight is being given to a perspective because the agent expressing it is powerful in the organization or particularly articulate, and etc.

Requirements

The broad aims of Practice Mapping are simply stated. We need to gather data sufficient to determine the behaviors of interest and the set of interacting agents to model. We need to produce practices that generate the behavior of interest. These practices need to be adaptive, and their adaptive interactions must reflect the underlying social/organizational structure (i.e., regularities in resource flows, communication, and role relations).

We are led to the following collection of requirements:

1) **Focus.** We require methods to determine what behaviors to focus on, since the phenomena we might attend to (the behaviors of potential interest) are infinite.
2) **Elicitation.** We need methods to elicit the range of participants, situations, action plans, and criteria that underly reports on behaviors of interest.
3) **Structure.** We need methods to collect and structure practices according to the three dimensions of social structure – resource flows, communication flows, and role relations.
4) **Validation.** Finally, these methods need to operate within a validation discipline.

We summarize our responses to these requirements in the following sections.

Elements of the Method

Focusing Methods

Mappers employ the concept of a **theme** as a primary focusing device in mapping. In a novel or stage play, themes are the framework and the moral of the story being told – they summarize the content as well as a value judgement on the content.[6] Similarly, in PM themes are action patterns (experienced behavior) that client-subjects recognize as having an influential impact on behaviors that they experience. As an example, clients often point to recurring arguments that have never been resolved and that harm productivity, or to significant kinds of situations that the organization never seems to respond effectively to, or that certain types of teams seem to do really well with a certain class of tasks.

In recent work on another project, we are developing a theory of human personality, based on the idea that personality characters (e.g., stubborn, dogmatic, true believer) are control regimes that describe large or important parts of an agent's practices. These personality characters are a kind of theme, at the individual level. In the case of asymmetric threats, we have built maps of agents whose behaviors exhibit personality characters within a wide variety of practices. For example, one apparently dogmatic agent divides his practices into ones that reflect one pattern of beliefs (e.g., the Islamist revolution) and others that embody beliefs incompatible with the first ones (e.g., partner with agents that sell drugs to Muslims).

We seek both positive and negative themes – the things that client-subjects prize, and the things that they believe are problems. Not surprisingly, we usually receive many more negative than positive themes.

Our protocols formalize the concept of a theme's impact in a common-sense way, namely, in terms of a **significance** component and a **frequency** component. For some applications, we have defined **impact** as the product of significance and frequency (i.e., a kind of expected utility estimate). More generally, themes are behavior patterns that client-subjects judge as very significant and more frequent than rare, or, if only moderately significant, then quite frequent. Equivalently, themes rule out all one-of-a-kind events, no matter how significant, and trivial events, no matter how frequent.

In our experience, client-subjects find themes very natural to state, and, they tend to recognize and understand the themes proposed by others – client-subjects do not tend to view themes as a novel or technical concept. We have developed protocols to elicit themes, as well as an instrument to enable client-subjects to rank them for frequency and significance. The resulting data offers many useful entry points to focus PM. For instance, mean and variance on the product of significance and frequency reveal important behaviors as well as conflicting views, both of which focus mapping.

Mapping can be focused by any other aspect of our work with an organization, not just by themes. For example, it can be focused by gaps discovered in the transition from mapping to modeling. It is also focused by validation methods, such as the exploration of conflicting views (e.g., to resolve conflict by uncovering differing experience in how practices are situated). Mapping is also focused to support intervention, such as technology insertion, training in new

[6] The theme "don't cry wolf", in the children's fable, both tells about a behavior and offers a value-laden judgement.

skills, or large-scale change management programs. But in general, extensive testing reveals themes as efficient and effective entry points.

Finally, mapping responds foremost to any client-subject goals associated with the organization deciding to undertake the process of mapping, modeling, and analyzing practices. Despite our care in developing Practice Mapping as a methodology, judgement is nevertheless required of mappers so that the potentially infinite amount of detail in the organization's practices is explored in a way that best serves organizational aims and decisions.

Initial Mapping: Variation Strategies to Map the Practices that Generate Themes

We now turn to mapping of individual and joint practices. As required by the definition of practice, mappers elicit information necessary to develop the action schemata that describe practices. Mappers seek to determine the situations that trigger a behavior, influence its execution, and result from its execution. Mappers also elicit the embodied action plan. And, mappers seek to develop the underlying criteria by determining the range over which behavior varies and then mapping the relationship between information that signals the need to adapt and the resultant behavioral adaptation.

This general requirement – to map situations, action plans, and criteria – applies at any level of practice. We begin with **variational strategies** intended to develop a preliminary specification of the practices that generate themes; this is the topic of this section. In the next section we describe a second layer of variation protocols, used to refine the maps first produced here into their specific resource, communications, and role-relations structure.

We begin by asking client-subjects to describe and illustrate in some detail the practices exposed by the themes they have produced. As much as possible, we seek here to allow them to describe it in their own terms. From this discussion, we have a preliminary understanding of the participants, stakeholders (perhaps), situational context, and basic action plan and typical outcomes. Also, because themes (by definition) are significant for the agents, we learn something about the values or criteria that are present or absent.

We then seek to develop a model of the practice(s) that generate the theme. Above all else, we tend to view every practice as a joint practice, and so we see agents as participants in an overall activity (e.g., participants in a meeting). We therefore seek to map the situations, actions, and criteria of the joint practice first, and then secondarily we seek to map the ways individuals contribute to the joint action. This discipline – joint practice mapping first, individual contributory practice mapping second – is necessary to capture the larger social criteria, and happens to offer an effective way to elucidate the three dimensions of social structure. This discipline also permits mapping to conclude at the level of the joint practice; in many applications this is sufficient.

Mappers now have the option of exploring the details of the practice, as well as the option of investigating the implications of the practice for other agents. In the former case mappers investigate the practices of **participants**. In the latter case, mappers investigate the practices of **stakeholders** – agents whose practices are influenced in their triggering or in their enactment by

the occurrence and/or products of the preceding practice. The choice of option can be evaluated according to the degree to which client-subjects are prepared for each discussion, and the degree to which the mapper seeks to understand a theme by exploring all of the practices that a) compose it and b) are influenced by it in the form of stakeholder relations.

Let us suppose that detailed mapping follows. Client-subjects will then be helped to articulate and come to consensus on the agents required, permitted, and prevented from participating in the behavior. Next, they will be helped to develop a shared view of how it usually unfolds or manifests itself, and the way this may vary under varying circumstances. Mappers employ a variety of variational strategies to develop the range of constituent situations, action plans, and criteria that are embodied. As just one of a large number of variational methods, we ask for illustrations of terrible, poor, acceptable, good, and excellent outcomes of the behavior, and then of the specific procedures and adaptations that would be necessary or sufficient to generate each type of outcome. This offers a wealth of information about the participants, actions, and criteria. As another example, we investigate the situations that seem to occur just before, during, and just after the various behaviors, and in this way build the situational context associated with the practice – as well as the special class of situations used to monitor and then guide adaptation. And as still another example, mappers propose variations and seek unacceptable or out-of-bounds versions of the action plan, to learn more about criteria. Note in all cases the criteria (or the values) suggested by the theme itself are guiding mapping.

Alternatively, mappers may seek to understand the stakeholder relationships – the agents/practices that are significantly influenced by the occurrence or particular form or products of a given practice (occurring previously or concurrently). Here, mappers seek to understand the preceding practices and the stakeholder practices as another joint practice, and so to map its situational and criterial and action structure accordingly. Note that the very nature of a theme introduces a boundary that has placed stakeholder practices on the outside. Therefore, this will be a 'redrawing' of the practice boundaries – at least, as commonly understood by the client-subjects. We find this shift in perspective is quite revelatory. The result of this mapping exercise is another joint practice description, amenable to detailed analysis and/or analysis of its downstream stakeholders.

Note also that new client-subjects will often need to participate in the mapping process. When mappers learn of new participants or highly-affected stakeholders, they have discovered new client-subjects whose testimony will help inform and validate the map. Further, mappers may seek to link the practices that generate themes together (i.e., so it is a fully-connected system of practices). Expanding out via stakeholder relations is a very effective way to accomplish this. In both of these ways, exploration of stakeholder relations makes the determination of the appropriate 'frontier' of considered practices and agents internal to the method.

We elicit not only experienced behavior but also imagined behavior. Agents have practices that they are aware of and disposed to execute, but have not yet executed (e.g., how they would react to emergencies they are aware of but have not experienced).

Refinement Mapping: Variational Strategies to Establish Social Structure

At this point, mappers and client-subjects have produced a map of the practices that generate the themes. Usually, they will have investigated stakeholder relations sufficiently to connect the (practices that generate the) themes. The map will also contain significant detail on the agents, general situations, courses of events, and criteria that participate in, pertain to or guide each practice. At this point mappers work with client-subjects to develop the detailed social structure. In particular, mappers seek to refine practice models – i.e., refine the situations, action plans, and criteria/monitors for the existing practices – by careful attention to the way each is influenced by variations in resources, interpretive and communicative processes, and role or authority relations.

Each dimension of social structure is examined by varying a **reference item** about values determined to be nominal for the practice, thereby producing data to describe whether and how behavior varies. For role relations, mappers generalize event trajectories into task descriptions, and then vary task and features of tasks to explore role relations. For communications relations, mappers use role-play or other methods to gather sample conversations, generalize their content as an item we call a topic, and then vary features of topics to study communication. To explore resource relations, mappers develop a nominal resource profile for a practice (physical and information constituents) and then vary within an associated reference space.

The asymmetric threat domain offers ready examples of each dimension of social structure. For example, terrorist organizations have the following main roles: leader, ideologue, moral authority, operator, and sponsor. Individuals assume one or more of these roles and manage them – perform their tasks – in relation to others playing other roles. Communication is about a collection of topics, some distinctive (e.g., Jihad); interpretation is managed by employing selected distribution mechanisms (e.g., interview, Fatwa). Resources include the funds and information needed to sustain and defend the organization and provide for its specific activities; some (e.g., particular pieces of information) are distinctive and reveal plans and inter-organizational relations.

Mappers and client-subjects can begin their analysis of the social structure embodied by a joint practice by investigating along any of the three dimensions. However, communication and resources tend to be managed in role-specific ways, hence, mappers tend to begin by mapping role relations.

Mapping Role Relations

In essence, roles are expectations agents have for the tasks that other agents will undertake, the way tasks are coordinated across agents (relations over complementary roles), and for the kinds of results that participant agents will tend to produce.

Our strategy for developing role relations has several steps, which we summarize here by way of introduction. The steps below should be viewed as a logical ordering of analysis, not as a fixed investigation sequence: later steps mostly formalize results of prior ones, but may also force return to earlier steps.

1) Background development of roles. This is a top-down or outside-to-inside step in the development of roles.

a) Review secondary sources – popular material, as well as cultural and historical studies – for the roles that typify the organizational field being studied. For virtually any type of organization, and strongly so in the case of asymmetric threat, there is a wealth of material on roles provided for by the culture and/or organization type. There is even more material on how specific agents have played specific roles.

b) Develop tests, to be checked later, for whether the socially-constituted roles discovered in this step are in fact present in the subject organization, and if so, how they are played. E.g., predictions for tasks that will be undertaken, by given agents, in given situations.

2) Develop the particular roles of the subject organization. In addition to socially-constituted roles, organizations have their own specific ways of 'dividing up the labor'. The ones that matter and that recur are the particular roles in the subject organization. (We refer to significant and recurrent patterns as themes.) This is a bottom-up or inside-to-outside step in the development of roles.

   a) Review data and if possible conduct interviews to gather data on significant and recurrent behaviors: the situations that precede and trigger them and the plan of action that is followed. Significance will tend to be measured by the goals of the mapping exercise (e.g., situations, participants, activities of special import).

   b) It is useful to begin the nominal or typical way individual and joint behavior tends to go, and then gradually work out to the more rare, and to the proscribed, cases. In any case, it is important to work out from the nominal cases, since this is required to place bounds on the practices that generate the behaviors.

   c) Describe the participants, their individual actions, and the stream of results that is produced.

   d) Cluster behaviors with a family resemblance into themes (e.g., those that exhibit a particular decision strategy operating over a range of similar situations).

   e) If possible, interview client/subjects to elicit behaviors that seem important. Gather multiple instances of each, and cluster as candidate themes. Present/test the candidates, then ask for other themes and behavioral illustrations.

3) Develop practices that are sufficient to generate the themes, in the light of socially-constituted roles. In essence, events and behavioral data are analyzed into the tasks that behavior seems to serve. Task descriptions are generalized to form practices (the generators of behavior) associated with participant agents. Roles are the total set of tasks of an agent, and role relations are the properties of complementarity among the roles.

   a) Review the behavior clusters (themes) just identified, and invent task descriptions that 'explain' the behavior. Tasks are given by the objectives their actions serve, and the constraints they place on themselves and/or (tacitly) agree to live within.

   b) Array the tasks within and between participants, to establish coordination relations within and between agents.

   c) Pay special attention to the way tasks are shared. The best way to do this is to design a task description that covers the tasks of all participants – the overarching objectives and constraints, and attendant methods. This view offers a framework for tasks in relation to others. The overarching task description will place constraints on the individual roles in the hypothesized role relation.

d) Develop the models of agent practices – situation, criteria, action-plan descriptions – sufficient to recognize and carry out the tasks, in the situations associated with the tasks (i.e., the antecedent and triggering events that make the task relevant). Infer criteria from whatever is preserved in the face of variation in the way the task is performed, and from the tendencies exhibited in the observed behavior (e.g., ways action is corrected, conditions under which it terminates).

e) As needed, revise practice descriptions so that the practices properly generate the pattern of interaction of the agents.

4) Critique the results of previous steps.

a) Check consistency of the socially-constituted roles with the particular roles we have discovered (i.e., the themes). In particular, check that the relationships between roles found in the larger social sphere are compatible with task relationships found in the organization. If not, then both need to be revisited. The real practice of the organization is the final arbiter; socially-given roles are just guides.

b) Seek counter-examples from client-subjects and secondary sources – violations of the role specifications.

c) As roles and role relations begin to be well described, mappers investigate how each participant agent perceives its role in the hypothesized role relation. For example, the tasks of some agents may be used to infer complementary tasks that are probably undertaken by other agents; in this case, it is necessary to test the hypothesized complementary tasks.

We now describe these steps in somewhat greater detail.

Client-subjects live their practices, meaning that tasks – objectives and constraints – are mostly tacit. Mappers begin the articulation process by simply asking client-subjects for 'how things tend to go'. From this mappers infer the typical execution of the joint practices that produce themes. Mappers focus on the sequence and inter-relations of central events that seem to occur. These events are the effects, in the world and on agents, both in process and as final outcomes, of the joint practice.

Socially-constituted roles should be investigated and then used as reference points in understanding the particularity of a given organization's roles. The institutions in the larger society determine the main roles and role relations, and how they are recapitulated in time by different agents – how agents are 'trained into' roles. Information on these roles and the training processes can be extremely informative when we critique our understanding of the particular roles we believe we are seeing in the subject organization.

Once the nominal event trajectory has been established, mappers and client-subjects establish the boundary of associated events. They do so by identifying unusual or rare trajectories, and as much as possible, the rationale that makes the alternative unusual or rare. Mappers seek alternative trajectories by proposing variations at the level of key events, or differences in the agents that produce the events – e.g., "can this practice go another way?". Mappers may also suggest alternatives – e.g., "in some organizations it goes this way; do you find this?", or,

"would this be allowed?". The boundary is established once we know what is ruled out for a practice, and why, e.g., "why can't it go this way?".

Data on the frequency of the various trajectories is maintained. Later, when practices are modeled, this information informs statistical distributions that are sampled during simulation studies.

Outcomes and certain kinds of constraints mentioned in this data are collected and used to suggest objectives. (Note that most practices have outcomes in the form of particular products, like decisions about a matter, and social functions, like perpetuating authority relations.) Information on trajectories and trajectory violations are used to suggest constraints. Mappers maintain records on which agents do what, under what conditions.

When ready, mappers generalize the objectives and constraints they have received by hypothesizing a task description. A **task** is an objective coupled to a set of constraints, and identified with a particular agent or agents. Objectives include goals and maintenance conditions. Constraints apply to the agents that may or may not undertake the task, the legitimate methods, any restrictions on the situation or available resources, outcome requirements, and etc. Constraints can be absolute (hard), or can express tradeoffs and preferences (soft).

The candidate task summarizes the joint practice's function, constraints that have been revealed, and the way constituent agents tend to participate. Mappers then seek critique by asking client-subjects whether the proposed task description fits their experience. The most effective way to do this tends to be to use the description to stimulate client-subjects to present counter-examples. Event trajectories and task hypotheses are revised as needed.

Roles are collections of related tasks that serve larger objectives. The inter-dependence of tasks, especially objectives, is used to assemble tasks into roles.

Some tasks are more important or central to the organization than others. It is useful to determine these first, and to make them references for exploring related and dependent tasks. Then, the roles associated with the key tasks are references to guide exploration of other roles in the organization.

Earlier, we said that roles are found in the expectations agents have for the performances of themselves and others. We observe here that the review of background material, and the analysis of joint behavior trajectories, has relied on mining expectations. Tasks are formalizations of expectations – summaries of expectations for particular performances in terms of the objectives, constraints, and methods that are specialized and applied in given situations.

A role relation is exhibited in the pattern of interaction found in joint behavior. Mappers study what each participant does in the larger joint action, how participants jointly assess the progress toward objectives, how action is synchronized among participants, and etc. Role relations are the properties we develop to describe the relationships. In other words, they are complementarity properties of roles.

Note that in highly formalized organizations some roles will have names (e.g., managing partner, chief counsel, secretary, marketing division, industry leader, transportation sector), and so mappers can uncover the existence of roles by asking for the types of agents in the organization. However, there are usually tacit/unnamed roles that are stable and important enough to warrant analysis, so this strategy is not sufficient. In any case, mappers view these named roles as labels for possibly-important and distinct task complexes, that is, mappers seek to use them, in addition to the practices identified in themes, as entries into the acquisition of tasks.

This bootstrapping process helps client-subjects articulate their experience of joint actions as kinds of shared tasks. As part of this, subtasks will be identified, and constituent practices and agents will also emerge, and the mapping process continues (recursively) at this lower level.

As formalization proceeds, mappers keep structured records, which we summarize in the form of the following relations:

(theme , {joint practice})	the joint practices that constitute/generate a theme
(joint practice , {event traj.})	data on event trajectories associated with a joint practice
(joint practice , {agent})	the participant agents in a joint practice[7]
(task , objective , {constraint})	a task description
(joint practice , {task})	the tasks that summarize the joint practice
(role , {task})	a role, described as a set of tasks
(agent , role , {practice})	the practices by which a particular agent plays a role
(agent , {role})	the set of roles played by an agent
(joint practice , {role})	a role relation, described as sets of complementary roles

At every level, the mapper is keenly interested in the way joint tasks are shared – in the distribution of authority, responsibility, contribution, etc., that determines the shape of the practice. As roles and role relations begin to be well structured, mappers can work with individual participants to 'get their take' on each role they play and its direct/complementary roles, as well as on the role relation. E.g., it is very enlightening to interview doctors and patients (separately) to gain their perspective on the tasks associated with the doctor role, the patient role, and the doctor-patient complex. Inconsistent perspectives are mined for counter-examples and clarification.[8]

[7] Stakeholders are participants of other joint practices.
[8] Of course, the inconsistency is often the reality. Mapping just exposes it.

Mappers analyze these various perspectives for the criterial pattern we claim is most emblematic of social structure – the way individuals orient their behavior primarily via criteria that seek to maintain particular states of the joint activity. In this case, mappers analyze how the objectives that apply to individual role/task contributions relate to their tasks (we call this **adapting**), to the tasks of others (we call this **co-adapting**) and, especially, to their view of the overall task (we call this **harmonizing**).

Once mappers have moved from event streams to hypothesized tasks (and thus roles and role relations) mappers can conduct variational analysis at the level of the objectives or constraints that define the tasks. Variation at this level very frequently devolves to variations in the resources or the topics of conversation that guide the joint practice; we describe variational strategies for these cases below.

There is an alternative for task analysis that proves valuable in some role and role-relation mapping situations. This is to map role relations by mapping the decision structures and decision inter-dependencies within and among agents. In some ways this is dual to task analysis: decision structure is the choice structure within practices that accomplish tasks and fulfill roles. Variation here proceeds by eliciting expectations for what happens if given decisions are made (tasks are undertaken) by other agents, who takes over for who, what happens if decisions are made in novel ways (in terms of process, values, information), and etc. In an extended form of this protocol, mappers can investigate role learning by investigating how decision-making purview has changed during an agent's tenure with an organization. This alternative protocol is particularly useful when the organization seeks to redesign its practices.

Mapping Communication Relations

Communications mapping is similar to role mapping, in that mappers generally proceed from event descriptions (utterances in send-and-receive-and-send patterns) to a generalization in the form of hypothesized topics.

For communications relations, mappers seek to facilitate the development of understanding of how joint practices employ information transactions and interpret their meaning. Mappers begin by helping client-subjects articulate their view of who tends to say what, when, and to whom.

It is important to learn which agents tend to initiate which kinds of conversations. Prohibitions and 'missing conversations' are as important as actual conversations (often moreso, in our experience!).

Mappers review transcripts of conversations, communiques, e-mail, voice-mail, and etc. All are evidence of the use of language, to further larger ends, in particular situations. Role-play can also be used to develop data on conversations: mappers ask client-subjects to act out conversations as they pertain to specific practices occurring under the range of conditions associated with the practice. Official organization documents are reviewed (e.g., corporate handbooks, policy statements, billboard postings), as well as wall art and any other communicative object.

We next investigate how perspectives lead individuals to particular speech acts – so that we learn the linkage between speech situations, initiating speech acts, receipt and interpretation, and then replying speech acts, for the various participants. Mappers gather this data in order to develop what we call a 'standard reading' that each participant has for these utterances – the participant's way of understanding the utterance in terms of actions (especially speech acts) that become relevant and/or proscribed.

Mappers develop and hypothesize topics from these standard readings. Topics are generalizations of the various speech acts into a summarizing description (e.g., so, this kind discussion was about how to limit the creation of unnecessary bureaucracy). Usually, mappers propose topics as candidate generalizations, and client-subjects test them and offer counter-examples as relevant. (In some cases it is possible to move directly to this step, by asking client-subjects for topics that are typical of communication within a joint practice.)

Almost universally, there will be multiple topics associated with a practice, and frequently different client-subjects will perceive or focus on different sets of topics. The diversity of topics reveals the perspectives that client-subjects bring to conversation (e.g., one sees the meeting as reporting on status, another as developing tasks to take away from the meeting).

Mappers then refine practice descriptions by mining topics for the items listed below.

- the data that client-subjects are (evidently) focusing on, relative to the topic
- objectives served by each speech act
- presumptions and entailments in each speech act
- the flow of conversation as acts of interpretation lead to acts of speech
- the way final outcomes of the joint practice depend on these communication processes

In addition to exploring the topical diversity that is immediately available in client-subject reports, mappers can investigate variations along any of the dimensions implied by the previous dimensions. Just as with tasks in role relations, mappers are interested in nominal conversation (locale, situation, topic, participants) as well as less-common, rare, and proscribed conversations.

For example, mappers can investigate what happens if topics identified with one context are brought up in another. Similarly, they can investigate what happens if the habitual way of discussing a matter is challenged, or how discussion varies as discussants or locale varies (e.g., what becomes discussible and undiscussable). These latter protocols can be seen as a direct strategy of varying about topics-as-references (i.e., in addition to relying on the present-at-hand variability inherent in topics).

The result of this analysis is a map of the way joint practices are coordinated by the inter-change of interpreted information. When the maps (associated with coupled practices) are linked together, a map of the overall interpretive/communicative flow in the organization is produced. This yields important information about the nature and quality of distributed decision-making.[9]

[9] In one application, an organization was plagued by pernicious conflict. Communications analysis showed that there were critical breaks in the maps – conversations were missing, or

Mapping Resource Relations

To map resource relations, mappers and client-subjects investigate how participants and stakeholders adapt their individual and joint behavior if resources are varied. They do so by developing a **resource profile** that establishes dimensions (hopefully quantitative scales) for the quantity, quality, access relations, and timing relations that can form a specific resource endowment when an agent has a given task. Mappers work with client-subjects and analyze secondary data to determine how behavior (practice execution) varies with varying resource endowments.

For example, physical resources may vary from minimal to opulent, from impossible to acquire to readily accessible, from low-quality to high-quality, with one delivery schedule versus another, and so on. Similarly, mappers and client-subjects explore the effect of an informational resource being absent versus present, and with likelihood varying from low to certain. As part of this, nominal resource requirements for tasks/practices will be established. Ideally, likelihood distributions, for various points on the resource profile, will be developed.

As behavior varies over the resource profile, mappers seek to learn how situations, criteria, and action plans vary. For example, criteria are exposed by elucidating what is preserved in method and outcome, versus what is dropped or adapted, as resources are varied. We also explore the options of seeking more or better resources (e.g., gathering information) and transferring surpluses.

In some cases, mappers may begin with by mapping resource flows, and then seek to map the decisions that guide the flows. Logistical and supply chain analyses are usefully examined in this manner.

Concluding Remarks on Mapping Social Structure

The variational strategies we have developed for each dimension of social structure tend to yield information about all dimensions of social structure. The separate variational strategies are required for completeness, not because they offer independent evidence for each dimension of social structure. Hence for any variation-based inquiry on any single dimension of social structure, mappers expect to learn more about the details of the practice, who each client-subject believes counts as a participant and as a stakeholder, which elements are understood (and how well) by which participants and stakeholders, and ultimately, about the dynamic interplay that determines each of the dimensions of social structure.

Each category of requirement described previously – establishing foci, developing classes of practice constituents, eliciting social structure, validation – offers an entry point for mapping. It

---

prevented by the restrictions of defensive and/or strategic behavior. Role analysis then revealed the tasks that preserved the communication structure, and then agents that undertook those tasks. This led to the bottom-up discovery of opposing factions that could not make effective joint decisions – exhibited directly in the role relations guiding communication tasks.

is not necessary, for example, to wait for full mapping of a behavior prior to exploring validation questions, or following an explicit focusing exercise.

We find that data on any behavior potentially informs many other behaviors and elements of social structure. That is, one does not set out to discover, say, situations in isolation, then action plans, then criteria. Instead, a report on an aspect of behavior is analyzed into the full range of constraints it places on all (elements of) the co-adapting practices that generate and are influenced by the practices.

It is important to note that testimony on behavior is always through the eyes/interpretation of the respondent, and in relation to the mapper's questions and interpretations. The final generative system of practices will draw on, and will have to rationalize, varied and conflicting reports. Note that some conflicts may be irreducible at the level of interpretation, and yet the practices that they describe are able to inter-operate; therefore, rationalization is first and foremost at the level of behavior.

At any point, a practice can be viewed as the lowest unit of analysis.[10] At this point the situations, action plans, and criteria are identified and assigned across agents. Nothing prevents later mapping that refines the practice into constituent practices and/or agents in a role relation.

The previous properties of PM induce a strong requirement for tools that keep careful track of the data that supports each modeled element. Also, as a practical matter, mappers need a reminder technology that alerts them to the state of mapping of various practice sets (e.g., noting a stakeholder has been mentioned but for which there are no practices). We have developed methods, and are developing technology, for this purpose.

Validation: Standards and Methods

It is common in the domains of crisis and asymmetric threat management, and even in the field of organizational analysis, to find that claims about phenomena and/or actors ultimately rest on the characterization of the person making the claim as "an expert". Such claims may have face validity, but they do not have scientific validity. Our objective in this section is to present the scientific standards that have guided our work.

We orient our discussion of validity around the following three types of standards: concept, construct, and criterion.[11] Each of these standards measures our ability to characterize some content field or phenomena of interest.

- **Concept validity:** the degree to which the abstractions used to describe content are linked to other well-founded concepts.

---

[10] Recall that a practice can be a primitive activity, a composite activity, an individual agent, or a composite agent at any level of organization.

[11] We thank Dr. Gary Jackson for his suggestion of this division and for his helpful remarks on validity as sought in the psychological testing community.

- **Construct validity:** the degree to which concepts can be effectively measured in the content field of interest. In our case, the degree to which mappers and client-subjects can use our concepts to recognize and interpret behavior.
- **Criterion validity:** the degree to which the content field can be represented as law-like, using the concepts and implications drawn using the concepts. Prediction, post-diction, and retro-diction are each tests of lawfulness.

It will be seen in the above that claims of validity are not final boolean judgements of correctness (i.e., something is either valid or invalid). Instead, *validity is a matter of whether a content characterization helps us produce useful claims (criterion validity) in an effective manner (construct validity) and integrates into other concepts that enjoy their own validity status (concept validity).*

We now offer a few remarks on each kind of validity as it applies to our research, and then present a few implications for practice mapping and technology development. Much more extensive discussion is available in other sources.

Generative Practice Theory is the content characterization theory we have developed. In terms of concept validity, our work draws upon, integrates with, and extends leading thinking the social and organizational sciences. In particular, we see our work as a constructive synthesis of the following main elements:

- Contingency Theories (cf., Galbraith)
- Institutional Theory (cf., Scott, Meyer & Rowan)
- Critical Theory (cf., Habermas)
- Group Dynamics & Action Research (cf., Lewin)
- Organizational Adaptation (cf., Giddens - *structuration*)
- Logic of Practice (cf., Bourdieu)
- Organizational Learning (cf., March, Wieck, Pfeffer)
- Network accounts (cf., Carley's Constructural Theory)

In addition, the action schema construct, and its algebraic and π-calculus formalizations, are contributions that draw on and extend the research in Computer Science on action representation. The agent architecture we have developed draws upon and extends leading proposals for resource-dependent real-time architectures in Artificial Intelligence.

We have engaged in an unusual amount of work with clients to test and refine our concepts and methods, and so have extensive information on construct validity. We mention a few cases. In city planning, the concepts of practices, themes, and mutually-generating dilemmas proved very fruitful. The model of organizational structure (resource flows, dialogue, role relations) led to the diagnosis of what had been a stalled decision process for the city. In our work with the EPA, we were able to produce a description of their practices for a) forming crisis response organizations, and b) responding to crises. The client/subject we worked with was the leader for U.S. West Coast development of community crisis response plans; he claimed that our model was both correct and superior to the one he had spent two years developing. In asymmetric threat studies, we find that practicing intelligence analysts make quick and effective use of our concepts. We

worked with nuclear proliferation experts, special operations and terrorism experts, and the FBI, and each time produced models of content fields that surprised and enlightened them. We have also been able to train junior intelligence practitioners in our method of Practice Mapping, with powerful results. Finally, we trained teams of students at the Army's Command and General Staff College in a framework for constructive, critical dialogue (in a few hours). In experiments, their crisis action plans were judged uniformly superior to those teams that did not receive the training.

Construct validity is very important for us, since we aim to endow organizations with a language (GPT), methods (Practice Mapping, Planning by Analysis, Critical Dialogue), and technology (ACCORD, CFA) and in so doing make them more effective. *We have proven that our work can be understood and used effectively by others, and so enjoys substantial construct validity.*

Criterion validity is a test of whether a characterization can produce a description of content. In other words, does the content field obey some laws, and do we have an approximation of those laws that we can use to produce a description of interest? We present the cases we have looked at below.

As a preliminary, we use the notation t^{--}, t^{-}, t^{now}, t^{+} to refer to ordered time intervals. Then **prediction** is making a claim, at t^{now}, about content at t^{+}. **Postdiction** is making a claim, at t^{now}, about content that was not but that could be found at t^{-}. It is a 'reverse prediction', or what is sometimes called a back formation. **Retrodiction** is making a claim, at t^{now}, that if one was placed back at t^{--} and armed only with what could be known then, that one could predict something about t^{-}. Retrodiction functions mainly as an explanatory mechanism (e.g., a famous example: given this pattern of tracks in the snow, we retrodict that if we went back in time and watched, we'd see a cat cross here).

We have two kinds of predictive tests. First, based on our theory, we have been asked by experts to offer an opinion on whether a particular event, for a particular asymmetric threat agent, would likely come to pass. The presumptions of the events violated our model of the agent, and we predicted (correctly) that the event would not occur. Second, we have diagnosed and intervened (successfully) to treat a number of organizational problems (e.g., agents A and B stimulate each other's defenses in ways that lock pernicious cycles in place and that seem then to ratify their limited and incorrect theories of the other). In these cases we have been able to predict recurring behavior in given situations, as well as design alternative practices and predict their outcomes once implemented.

In general, intervention always involves prediction along these lines – and so to be able to successfully predict and alter a system is evidence of the validity of the model of that system. Moreover, intervention inherently produces evidence with which to test and extend the model of the system. Our work in city government, community dispute resolution, doctor-patient relations, and crisis response planning exercises each successfully exemplifies predictive criterial validity.

We have one test that can function predictively, postdictively, and especially, retrodictively. We conducted what is known as an intellective study of an organizational learning case (on another contract). Intellective studies hypothesize some relationship – in our case, the amount of learning

that would tend to occur under varying conditions of conflict and time pressure for results. We hypothesized that learning would be low in low-pressure situations (little conflict, little time pressure) and in high-pressure situations (large conflict, significant time pressure), and that learning would increase monitonically to some maximum between these two endpoints. These studies then simulate the variation in the content field and test whether the hypothesized relationship is found. If so, then the relationship can be sought in real data. Typically, this will be a retro-diction used to explain past behavior, but it could be a prediction or even a postdiction. Our results clarify the literature on organizational conflict and groupthink, and offer clear guidance for what to measure.

We have an important case of retrodiction in our characterization of South Africa's development and dissolution of its nuclear weapons program. We were able to show that, empowered with our theory, analysts would be led to look at different data sources (e.g., socio-cultural material on the Broderbund), and would interpret and weigh all data differently than was actually the case. Doing so leads directly to the prediction that South Africa would have a nuclear device at least two years before they did, and so (depending on who you consult) two to five years prior to when the U.S. decided they did. The proliferation and South Africa experts we worked with now see the case in the light of our view.

We conclude our remarks on validity with a few observations on methods, and technology support tools, that support the creation and testing of claims (and thus, the validity standards our claims might meet). Whereas concept, construct, and criterial validity aim at large-scale tests of a theory and its methods, the techniques we mention now are about the fine-grained tests, record-keeping, and low-level methods that implement the large-scale tests. These methods help the theory and its tests function over the largest possible content range.

Every mapping protocol must, as a zeroth-order validation method, seek not only support but also the possible refutation of reports on behavior. After all, client-subject reports on behavior are always limited and sometimes false: articulation is a practice and so is limited by our capacities, experiences, and the perspective we bring to the practice. To accomplish this, comparison with observable behavior and its results should be made whenever possible. Testimony should be analyzed for their embodied perspectives, and when possible confronted with alternative perspectives revealed when others are interviewed and/or are invented by the mapper – in this way articulation, as a critical process, is a 'bootstrap method' for producing results with greater clarity and validity.

We test for inter-subjective agreement by client-subjects that the map correctly depicts/generates their descriptions of practices. This usually occurs by confronting client-subjects with the diverse and conflicting views of others and then facilitating a refinement process that creates a consensus view (e.g., adding contingent structure to practices so that the different views reflect different situations or cases).

We also test whether client-subjects find the map not only captures their experience, but also offers important insights for understanding that experience. It is on the basis of new and useful insights being produced that our maps have compelled many of our client-subjects.

Our technology supports these methods. In ACCORD we maintain links, for each element in an organization model, to the data on which it is based. We explicitly record data that supports, as well as data that refutes or weakens belief in any structure in the system. To do so we have developed the capability to depict ACCORD structures as claims of full-fledged argument structures that depict data and the reasons that determine whether they support or weaken belief in a conclusion, and how they do so. Hypotheses and assumptions are similarly maintained. An evidential reasoning calculus could be added.

Similarly, the CFA is required to maintain arguments in support of organization models. Critiques are dialogue transcripts of argumentation processes, and are maintained as part of the rationale for each model and its evolution. This is essential support for a collection of analysts and decision-makers using these technologies.

Both ACCORD and the CFA maintain anomalies as first-class objects. This is a particularly important type of critique, and is the one that central to Planning by Analysis. Our growing experience with organizations teaches us that anomalies are in fact 'sprinkled out' over time and throughout the organization. Therefore, collecting them produces an invaluable resource: what for any agent might just be a local oddity may become, when combined with the anomalies recognized by others, the clue to a much deeper understanding. Since human systems adapt and thus are always evolving, this is type of support appears to us to be essential.

Planning by Analysis

Armed with GPT and Practice Mapping, modelers, planners, and decision-makers have the necessary data to build formal computational models of organizational crisis response, or for that matter, any psycho-social dynamic of interest. Yet the primary value for organizational decision-making of any kind lies not in the models per se, but rather in the insight they offer into phenomena, and the support they provide for taking more effective action. Simply put, modeling is in service to decision-making; it is rarely an end in itself.

We have invented a discipline that places mapping and modeling within a larger problem-solving framework that we name Planning by Analysis. **Planning by Analysis (PBA)** is a process that combines problem-solving tasks (e.g., planning, deciding, communicating, acting, and action/plan monitoring), with the organizational mapping, modeling, and analysis actions. The former generates action plans, whereas the latter produces and tests the understanding that informs these action plans. In this way, PBA seeks to balance responding to events in a timely way with acting to improve the understanding on which response is based.

What's Planning in Planning by Analysis?

The generative theory treats almost all action as joint action. Hence the mapping and modeling in PBA entails taking the distinctive step of modeling oneself as well as all other participants and stakeholders. Planning by Analysis is a way of exploring the results of self-other behavioral prospects, and of refining both the models of self and other, and changing the action plans that guide self, until effective practice systems are designed.

In its simplest form, PBA begins with building a model of the social system one is interested in: usually, one's own organization and its coupling relations to the organization (or social system) that forms the interactive context of one's own actions. For example, to plan a non-combatant evacuation operation in a foreign country we might describe the adaptive practices of insurgents, government officials and the troops of the foreign country, our embassy and its staff, our evacuation troops, and any other agents of importance. Simulation experiments are then designed, executed, and analyzed under various scenarios that the modeler thinks relevant for behavior of other agents in the world and/or for self actions.

Simulation experiments project implications of the agent models under the varying scenarios. In the process, these experiments produce a range of consequences and estimates of their frequency of occurrence – each of which may be subject to empirical test (i.e., "do these things happen?", "this frequently?"). As such, PBA experiments inform the experimenter not only about the potential futures of the agents of interest (prediction) but also of the quality of the agent models (validation) and even the quality of the explanation of past agent behavior (retrodiction).

Plans emerge as superior options in the light of important scenarios and validated agent models. In other words, a plan is not just a partially-ordered set of steps that achieves a stipulated goal state from initial conditions. Rather, it is a product whose development includes the careful articulation and validation of the social dynamic within which the plan will execute. Further, the simulation studies can be used to develop expectations that in turn establish how best to monitor plan execution.[12]

We claim that this form of planning is much more powerful and appropriate for crisis response in the emerging world order. For example, asymmetric conflicts and so-called operations other than war have a vastly more complex and influential social dynamic than experienced in major force-on-force engagements (e.g., WWII).

What Guides the Steps in Planning by Analysis?

Individuals and organizations are frequently compelled to take some action. Rationally, we want to take the best possible action. At the same time, action is always informed by some level of understanding – of the situation, of one's potential or capacity for actions of various types, of likely results, of one's aims, of the capacities and aims of others, and etc. As a basis for the best possible action, we wish to have the best achievable understanding. But human boundedness, particularly in crisis or high-complexity situations, implies that we cannot always have the best of both – we cannot have the best possible plan, and the best understanding, unless we can afford to arbitrarily defer acting.

---

[12] Planning by Analysis can make effective use of the sorts of planners developed in the AI community, and in fact we see this as desirable. AI planners are silent on the plan vs. improve-understanding tradeoff, tend to lack or have very primitive models of other agents, and tend toward open-loop versus adaptive plans – deficits that PBA is designed to address. However, they are very useful in routine (if compute-intensive) searches through systems of physical and high-level behavioral constraints.

PBA is centered on the fundamental choice, or even dilemma, posed by these two wants: should an agent take its best action based on what it now understands, or should it postpone its action until it has taken steps to improve its understanding? More generally, agents are faced with the decision of how best to partition their resources between three activities:

- planning/deciding and managing action 'in the world' and based on the current understanding,
- acting to improve that understanding – which may delay planning and action in the world, thereby imposing the cost of delays but offering the benefit of a superior base for planning, and
- monitoring and evaluating whether the relative proportion of the other two activities should change.

Our response is to employ a bootstrapping strategy – begin at a reasonable point and then modify commitments in a sensible manner.[13] Agents start with a partition of the above three activities (acting, improving understanding, monitoring these two) that expresses default/current commitments. They also begin with whatever preparatory modeling has been undertaken. The proportion of effort devoted to planning, versus to improving the underlying understanding, is then varied according to the presence and substance of anomalies, as we now explain.

At all times, agents monitor for **anomalies**. Anomalies are significant discrepancies, instances where the implications of an agent's current understanding are silent on or at variance with trusted reports on world (physical, cognitive, or social) state, and where the difference is deemed to have significant consequences for things the agent values.

If the current understanding is incomplete or incorrect in substantive ways, anomalies will tend to appear, and their significance will tend to emerge if the discrepancies are about things that matter to the organization. In this way anomalies can be expected to 'pile up' and can be used to signal the need to rethink the guiding perspective. Then, in virtue of their specific content, anomalies offer ingredients (clues) for the development of alternative perspectives – ones that challenge current understanding but that better explain the data that is otherwise anomalous. Frequently, anomalies reveal assumptions that are tacitly embodied by the guiding perspective and that are now shown to be invalid.

The task of reviewing anomalies (to challenge the current guiding perspective) is a shift of attention away from acting and toward improving understanding. When anomalies lead to a new guiding perspective, information requirements will usually be suggested that both fill out the candidate perspective and offer tests of its validity – indicating a still-greater shift from a focus on action to a focus on improving understanding. In general, validation experiments should be conducted, via simulation and in the world, to test and affirm the new perspective. Models are updated as appropriate, and then plans are revised in the light of the updated models. At this

[13] In some cases there is no "beginning point" – the agent is always allocating effort among the tasks of acting versus improving understanding. Here, a bootstrapping approach seems mandated.

point, organizational effort and attention can shift back toward acting in the world, based on what is the new but now-current understanding.

In sum, anomalies are indications of the limitations of our models/understanding, and they signal when a greater proportion of the agent's effort needs to be devoted to improving understanding. This results in a change in the distribution of effort that increasingly favors data gathering (mapping), modeling, and analysis activities that if successful lead to improved models and then to improved plans and decisions. Planning by Analysis, by managing modeling effort according to the strength of existing anomalies, is thus an adaptive strategy for knowledge and action management.

Concluding Remarks on PBA

PBA stipulates that anomalies be collected stored for later review, for example, when they are or might be relevant to high-consequence decisions. The literature on crisis suggests to us that this is a particularly powerful tool for organizations, for the following reasons. Anomalies tend to be distributed over (discovered by) a variety of participants in the organization – individuals have different tasks and perspectives and so tend to be confronted with different anomalies. Also, our theoretical view and some empirical data suggests that we humans tend to 'explain away' local discrepancies. This is a product of a bounded reasoning heuristic of the form 'it is not necessary to change one's perspective every time something novel arises'. For both of these reasons organizations tend to fail to collect and then take advantage of the resource offered by the pool of anomalies. If collected, anomalies discovered by one agent can both offer additional critical tests of the views of other agents and, in their content, reveal important features of a better understanding.

PBA is a life-cycle conception for planning, that is, it views each element of understanding and action as both occurring within a stream of other events and evolving to adapt to those events. In light of this, data that is received through time and prior to the recognition of any given crisis should be encoded as mapped data or as fragments or constraints on GPT models. We refer to this as **preparatory modeling**; it is important both for the discipline it offers to processes of data capture and data storage and for the resource it offers, in the form of on-board model components, when crisis erupts or some other planning requirement arises. A fully-implemented version of our organizational analysis framework would also supply users with a library of templates for kinds of practices, agents, and organizations.

The PBA method can be expressed formally as a decision or risk management exercise. This could be carried out in two stages, a diagnostic stage and then a decision stage. First, the information about the world would be represented as a probability network (joint conditional distribution). When anomalies arise, variables are assigned the anomalous value for exploratory purposes (i.e., assume for what if purposes that the anomaly is valid), and the probabilities of other variables are computed. This generates a set of variables whose probabilities, if conditioned by the assumption that the anomalous data is true, seem unreasonable. Some of these may be easy to gather information on. The second stage then calculates the value of information associated with improving our information about one of the selected variables. To

do so it is necessary to describe the decision basis associated with actions that may resolve the anomaly, versus other options that encode the other primary alternative of continuing planning and acting based on what is currently known (i.e., without resolving the anomaly). Multiple sub-options in the form of different combinations of information-gathering, modeling, and model validation actions are associated with the main alternative of trying to resolve the anomaly and improve decision-making as a result.[14]

Critical Dialogue

Complex organizational activities (e.g., crisis response) generally demand the coordination of many agents and organizations. Its not enough that every agent do her part; actions, and the understanding of events and aims, have to be coordinated. In principle, any of the tasks of Planning by Analysis (planning, deciding, communicating, acting, monitoring, mapping, modeling, gathering information) may need to be coordinated. Similarly, the products of these tasks, and significant portions of the understanding and/or assumptions that inform these tasks, may need to be coordinated. **Critical Dialogue** is a form of structured communication that we have invented to support of multi-agent coordination.

Critical dialogue introduces two kinds of standards for quality communication. The first type of standard asks that communication be structured according to an argument form proposed by Toulmin. The second type of standard seeks to create a 'fair and open' forum for discussion. Together, these standards underwrite coordination by creating conditions in which participants are well informed, claims are well tested, commitments are the product of informed consent, and access to participation (i.e., to argument formation and evaluation) is fair and open.

Toulmin's argument form is as follows. The assertions of any agent are treated as claims and may have attached qualifying statements. Claims are depicted as based on or derived from evidence. A special structure, called the warrant, defines the relationship between evidence and claims (e.g., how the claim may be derived from the evidence). Another structure, called the backing, provides a rationale for the warrant.[15]

The second type of standard defines criteria for productive discussion; these standards ensure fair access to argumentation, and rational assessment of individual contributions. These standards are implemented as practices that seek to prevent power or authority, institutional arrangements,

---

[14] For some crisis or other situations, time might not be available to conduct this type of rigorous assessment of the central PBA decision. The formal approach is appropriate for prior/offline analysis; such analysis can be used to develop heuristics that can be used in real time (e.g., variables to which decisions/plans are very sensitive might be identified in advance).

[15] In a more general form we have produced called a Belief Structure, each element of an argument structure is the claim of a supporting argument structure (and so on, recursively). Belief Structures are composed of three interacting systems: an evidence system, a criterial system, and a support system. The support system produces assertions (claims) based on the evidence and the constraints imposed by the criterial system; it is responsible for maintaining the quality of assertions at all times (i.e., integrates derivation with truth maintenance). Recursively, evidence, criterial, and support systems are claims of more fundamental Belief Structures.

articulacy/persuasiveness or inarticulacy, and etc., from influencing access to or contribution to a conversation in a way that over- or under-values the evidential merits of the contribution.

In practice, individuals have varying skills with argument formation and even with particular aspects of argument formation. Mindful of this, we have proposed a collection of dialogue roles used to enhance the quality of argument formation and revision as a social process. In essence, a **dialogue role** is a practice of monitoring for problems with an argument or with the argument formation process (or opportunities to improve), and then critiquing the argument or the argumentation process accordingly. Hence these dialogue roles embody a way to satisfy both types of standards.

Sample dialogue roles include the following. The role of **Skeptic** is specially concerned with testing the quality of evidence and warrants. The role of **Devil's Advocate** is specially concerned with challenging assumptions (e.g., the scope or distinctions used for what is considered) and inventing alternative formulations for the available evidence. The **Gatekeeper** facilitates and manages access to information relevant to ongoing deliberations. The **Leader** manages participation (access to and influence on the discussion), and has authority to determine how much effort or attention a line of reasoning will be allowed. In this way Leaders shape the relative contribution of evidential considerations and participation, to produce the best possible result within the desired (egalitarian) forum for dialogue – all within the constraints imposed by the crisis.

The roles were developed based on an extensive review of transcripts and documentation on historically important crisis response situations (e.g., Cuban Missile Crisis). They were also evaluated experimentally, at the Army's Command and General Staff College.

In the experiments, expert planners at the Staff College presented students with crisis scenarios; students were to construct a response plan under fixed time limits. Half of the students received the standard training of the Staff College, and half additionally received limited (less than two hours) training in Critical Dialogue. The expert reviewers then evaluated results to see if the ones receiving the additional training: produced better plans, evaluated a greater range of evidence or options, evaluated options or evidence with greater care, or were able to produce superior justifications for their plans. The reviewers assigned much more favorable ratings along these dimensions to those participants that received training in Critical Dialogue than they did to those who did not receive the training.

Technology Framework: ACCORD Within the CFA

We discuss our technology development thrusts in this section. These are the **ACCORD** system for organizational analysis, and the **Consensus Facilitation Architecture** for collaborative work. A later section includes screen captures illustrating the use of the ACCORD technology to provide Planning-by-Analysis support during a hypothetical insurrection.

We have prototyped an environment that we call ACCORD, for "Agent Cognition and Coordination of Organizational Dynamics". ACCORD contains three suites of technologies for

conducting organizational analysis: mapping (data capture), modeling (data formalization as computational models) and analysis (visualization and analysis of simulation data).

ACCORD also includes a discrete-event simulator we have built to run simulation studies of organizational psycho-social dynamics. The simulation controller operates over the internet. In recent work, we are developing the necessary communication and simulation protocols so that simulations can be run over a (distributed) set of processor resources. The development of the simulator has occurred under other projects.

The ACCORD simulator is implemented in C++. This prototype will soon execute over a distributed network of Unix workstations, some of which are themselves multi-processors. The ACCORD mapping, modeling, and analysis technology suite runs in Java and C++ on laptops. The analysis suite includes the facilities to design and manage simulation experiments; it is able to communicate with the simulator over the internet.

We have also developed a preliminary specification of an integration environment for multi-user problem solving that we call the Consensus Facilitation Architecture (CFA). Among other things, the CFA reflects our view that, for any substantive organizational analysis, each of mapping, modeling, and analysis is likely to be conducted by groups of individuals. As such, there will be varying perspectives and even conflicts over data, assumptions, model elements, and analysis objectives. The CFA provides special support for the collaborative analysis and development of consensus over these various analysis constituents. The distributed discovery, analysis, and resolution of anomalies (as per the Planning by Analysis methodology) is a primary example of the use of the architecture.

The CFA is intended to run in C++ on multiple Unix workstations. Once implemented, ACCORD data structures will be sharable though the CFA.

We describe each of these technology developments in more detail below. We also describe a computational primitive, called a scheme, that we have developed and which we are using to implement the essential technologies in our program.

The Consensus Facilitation Architecture

In previous work for DARPA – as part of the DRPI Planning Initiative – we designed a system we called PS3. PS3 stands for "Problem-Solving Support System". PS3 is a software architecture intended to support distributed concurrent development of crisis response plans by automated planners. PS3 included the specification of two types of distributed truth maintenance capabilities that we invented to help maintain coherence over plans and plan rationales.

We reviewed PS3 to determine whether it could be used to support distributed collaborative human planning, augmented with automated problem-solving agents. We have generalized PS3's specification so that it now supports human collaborative organizational analysis. We refer to the new architecture as the Consensus Facilitation Architecture (CFA).

The CFA is a real-time, interruptible, distributed operating system that supports collaborative organizational analysis among distributed agents. It permits users to share environment data as well as mapping data, practice/agent/organization models, and conclusions or decisions. Further, we propose that a user's workspace in the CFA organizes data and models in terms of decisions and plans that are contemplated; after all, modeling is in service to action. In this way the CFA supports multiple inter-dependent mapping, modeling, and analysis activities, all in service to the framing, development, execution, and monitoring of coordinated decisions.

The CFA facilitates consensus (supports collaboration) by ensuring that transactions abide by the commitments of Critical Dialogue. In particular, the CFA calls for special facilities for users to a) form critiques of their products and others', b) share products and critiques over a selectable scope of users, and c) maintain rationales for models (i.e., treat each model component as a claim). Rationales maintain an argument in the form of evidence and warrants. Evidence can have the form of external data, modeling choices, and even modeling 'dialogues' in the form of multiple rounds of transactions that share critiques. Warrants represent the processes that convert ingredient evidence into specific model components.

The rationale or argument for each claim is maintained in the CFA, and these claims and rationales are shareable in support of consensus formation processes. It is therefore important to keep records of who happens to believe each claim or argument. With this in mind the CFA maintains two types of scoping rules: scope of publication, and scope of consensus. The first scoping rule marks claims according to the sets of agents who are permitted to have or who have had access to them – in other words, the scope of potential and actual publication of a claim.[16] Many claims and their supporting arguments will be marked either as private to the agent that formed it (e.g., work in progress) or public to all participants. The second scoping rule marks claims according to the set of agents for which the claim is treated as the consensus view – in other words, the scope of commitment by users to treat the claim as if is a truth (until refutations warrant reconsideration). By scanning claims and supporting arguments according to these two views, crisis responders (or software agents) can discover domains over which consensus is required, relative to the emerging response plan and/or the maps and models and analyses that embody the understanding on which the plan is based.

As mentioned previously, the CFA is intended to support coordinated decision-making. With this in mind, we argue that the CFA should permit data, claims, and rationales to be partitioned according to the way they support organizational tasks. Our current view is that four perspectives should suffice – a data, map (i.e., all ACCORD structures), plan, and milestone (i.e., plan monitoring) perspective. Each element within each perspective is linked to all layers of underlying rationale, and is marked via the scoping rules. The milestone perspective supports coordination in service to shared goals, so that users can measure whether their contributions will be produced in time, and reasoning backward from their needs, can make requests or impose constraints on the timing and products of others.

---

[16] A software agent could be developed to build the organizational scope of publication of a claim whenever multiple individuals publish it (e.g., A tells B and B tells X and Y).

The CFA specification will be enhanced to include an interface protocol for human users as well as software agents. CFA users will have access to the simulation engine and ACCORD as part of their tool suite.

The figures below present specific mechanisms to implement these commitments.

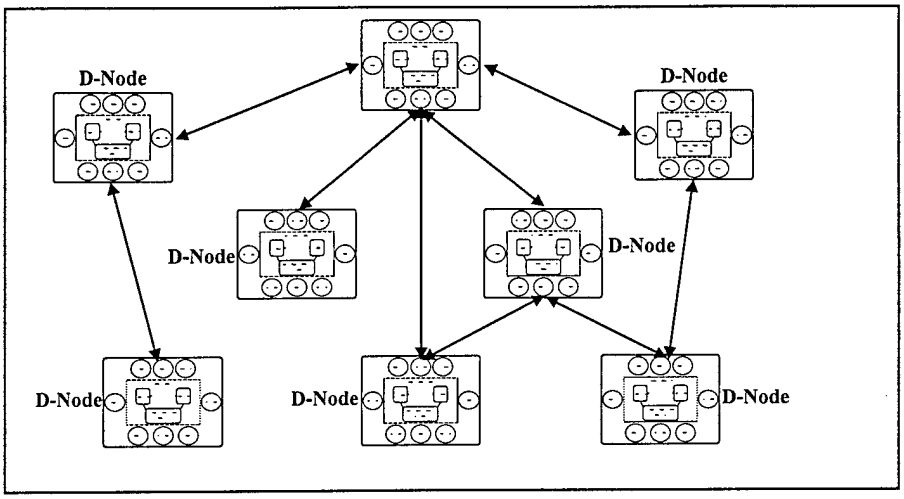

Here we show the "CFA Network Architecture"; it is a collection of two kinds of nodes, C-Nodes, and D-Nodes. D-Nodes are distributed problem-solving nodes of particular user agents. When D-Nodes point to other D-Nodes, it means they share data. C-Nodes are consensus formation nodes. When D-Nodes point to C-Nodes it means that the participants are publishing models with the intent of engaging in consensus formation about the models, its elements, its implications for planning and decision-making, or the rationales for any of these. C-Nodes add little functionality to D-Nodes; mainly, they make certain patterns of data visibility and sharing feasible and efficient.

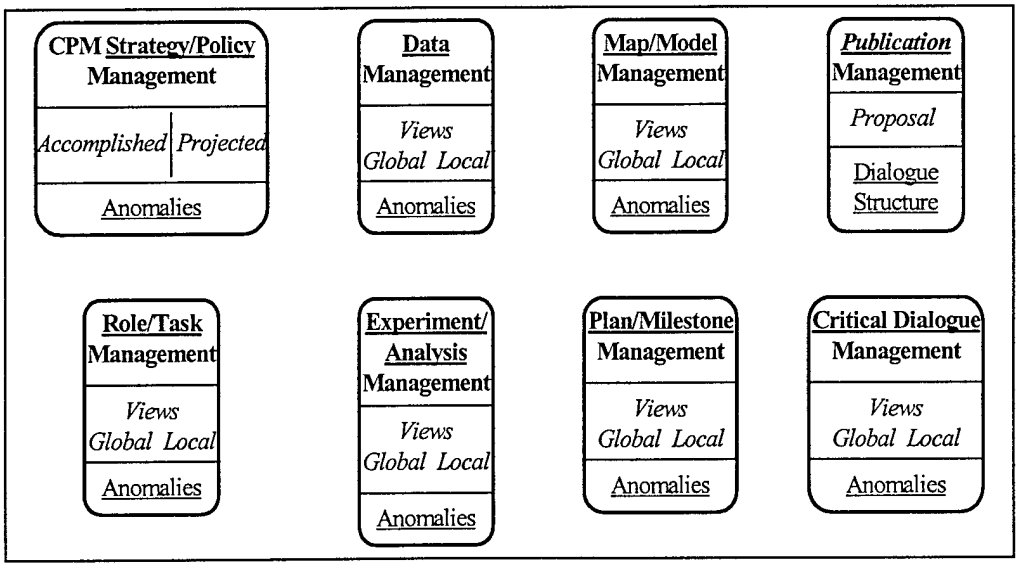

This figure defines the tasks or requirements we place on D-Nodes and C-Nodes. Our discussion of the CFA implies a need to store, manage, share, and revise data, models, plans and/or decisions, analyses, and progress toward milestones. All need their associated rationales to be available. Most importantly, users need to see how their local view (e.g., their part of the planning) fits into the global or consensus view. They also need to keep track of what they have accomplished, what is pending, and any anomalies in any of their work or the work of others that is visible to them.

Critical dialogue is managed by having C-Nodes optimize the sharing and presentation of the work of others as views at each D-Node. C-Node formation can also activate special tools that assist with consensus formation, such as intelligent software agents that review model fragments and their rationales to explore for points of overlap and inconsistency.

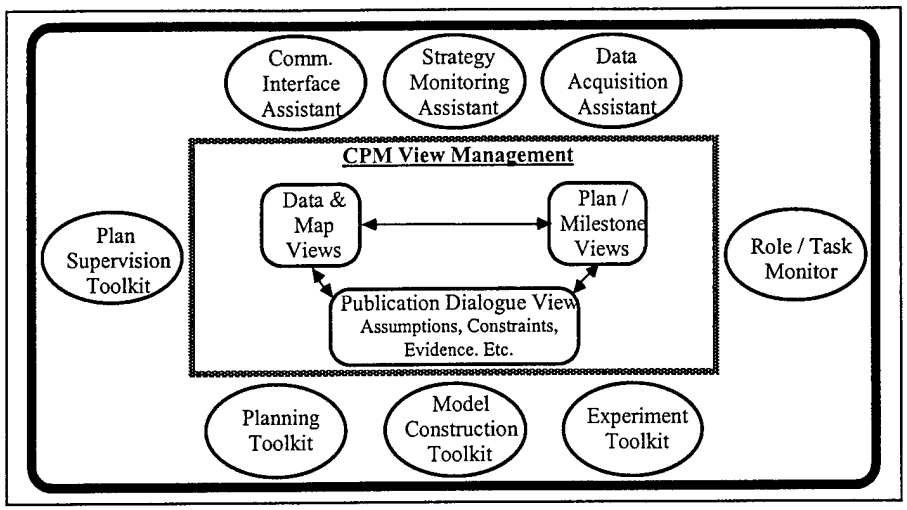

This figure depicts the D-Node architecture. In essence, it converts the requirements to support local work and consensus formation into a collection of modules that implement those services.

However, there are a few special features to note here. First, there is a collection of toolkits. These toolkits allow for the CFA to import any sort of tool natively employed by a user, such as an AI planner, or a decision analysis software package. Second, we show a collection of assistants – intelligent software agents that support complex data management and communication protocols, and that make available the status of the work of others. Third, we show a role/task monitor; its function is to monitor for new tasks (e.g., as revealed by products of the strategy monitoring assistant), and to advise the user on ways that current structures need to be made to fit into global structures. Hence this monitor assists with the detection of consensus formation opportunities or requirements. Finally, we show a view manager that organizes the data, map/model, and plan views, and that establishes the scoping on each.

ACCORD

ACCORD is the suite of mapping, modeling, and analysis tools that implement the distinctions within GPT and that support our view of organizational analysis under the Planning by Analysis discipline. We have prototyped and demonstrated the evolving ACCORD technology.

The mapping technology suite stores raw data on organizational practices, as gathered by mappers. It includes facilities to structure it in the form of agents, their action propositions, and (the three dimensions of) social structure.

ACCORD modeling technology supports the conversion of mapping data into practice, agent, and organization models suitable for computational execution. In essence, collections of action propositions are converted to computationally executable form by formalizing situation data as triggers and binding relations, criterial data as specific monitors, and procedure/action data as action plans.

The analysis technology supports the specification of monte-carlo simulation experiments – to develop implications of models for validation, information-development, and decision/planning purposes. This suite includes facilities to design scenarios over which monte-carlo simulation runs, to configure the experiment's stopping criteria, and to insert probes to measure simulation products during execution. The analysis tools also support visualization and analysis of simulation products. We are also gathering a variety of analysis tools that can be employed to evaluate mapping data directly.

Mapping, modeling, and analysis structures and results are maintained in a network of rationales that depict the basis of each item as a claim and/or a constituent of some other argument. Of these, perhaps the most distinctive and important one is the explicit support for users to store anomalies: the anomaly structure opposes real-world data and its basis with experimental results and supporting rationale in the form of data, models, and experiment designs. The stored anomalies thus provide the ingredients needed in Critical Dialogue and within Planning by Analysis to revise the organizational analysis (e.g., in service to adaptive crisis response) by critiquing and improving the understanding on which it is based.

We conclude this section with a collection of screen captures of the ACCORD technology prototype. These pictures show the use of the technology as illustrated in a hypothetical NEO (non-combatant evacuation operation). With the help of special operations, information operations, North Africa, and crisis response planning experts, we developed a high-fidelity scenario of an insurrection in Dar Es Salaam, Tanzania (about 6 months before the embassy bombing there).

When we demonstrate the ACCORD system, we show how models of the rebels, native government troops, embassy personnel, and U.S. troops, are developed and used to support evacuation planning and real-time evacuation operations management. It is an instance of what we call Planning by Analysis. The outline of the demonstration is as follows:

1) We begin with data gathering, modeling, and planning that occurs as the evacuation force steams to Tanzania from the Persian Gulf.

2) We inject events during the course of the scenario that are anomalous and show how users would capture the anomalies and use them eventually to reason forward to new data they should study.

3) A break occurs, and users review anomalies. This leads them to hypothesize the actor that could cause the anomalies is a cadre within the host-country troops, and conjecture there is a coup occurring under the cover of a general uprising.

4) Data is collected to fill out and test the new hypothesis. Agent models are revised (e.g., the government troops agent becomes two agents with different objectives: loyalist troops, and breakaway troops).

5) The nominal evacuation plan is now tested in the light of the new models and shown to have serious defects.

6) The plan is modified appropriately, and then satisfies casualty criteria for the evacuation.

7) Finally, the evacuation proceeds and is successful.

On the pages that follow we show pictures of selected points along the way in this scenario, which we use here to illustrate some of the technologies we have prototyped in the ACCORD technology suite.

This figure shows the top level control pane, allowing the user access to the mapping, modeling, and analysis suites. The labels Events and E-Mail are just stubs in this prototype; they are intended as ports offering to access the user's D-Node environment as well as other specialized data resources.

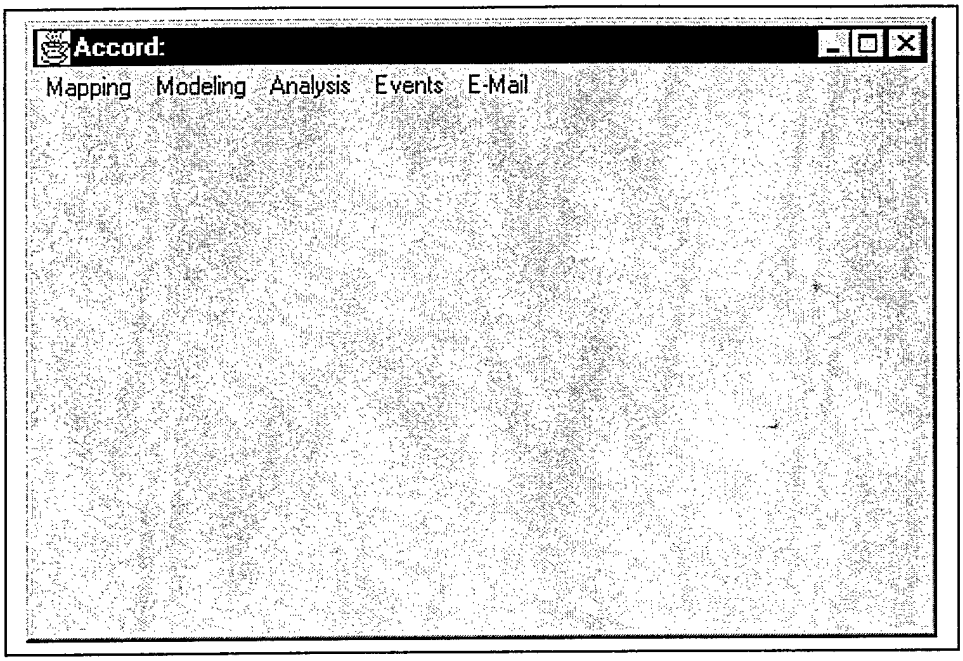

The figure below depicts some raw data, an extract from a recent dispatch of the ambassador. Note that we keep track of the source and anything critical of the source.

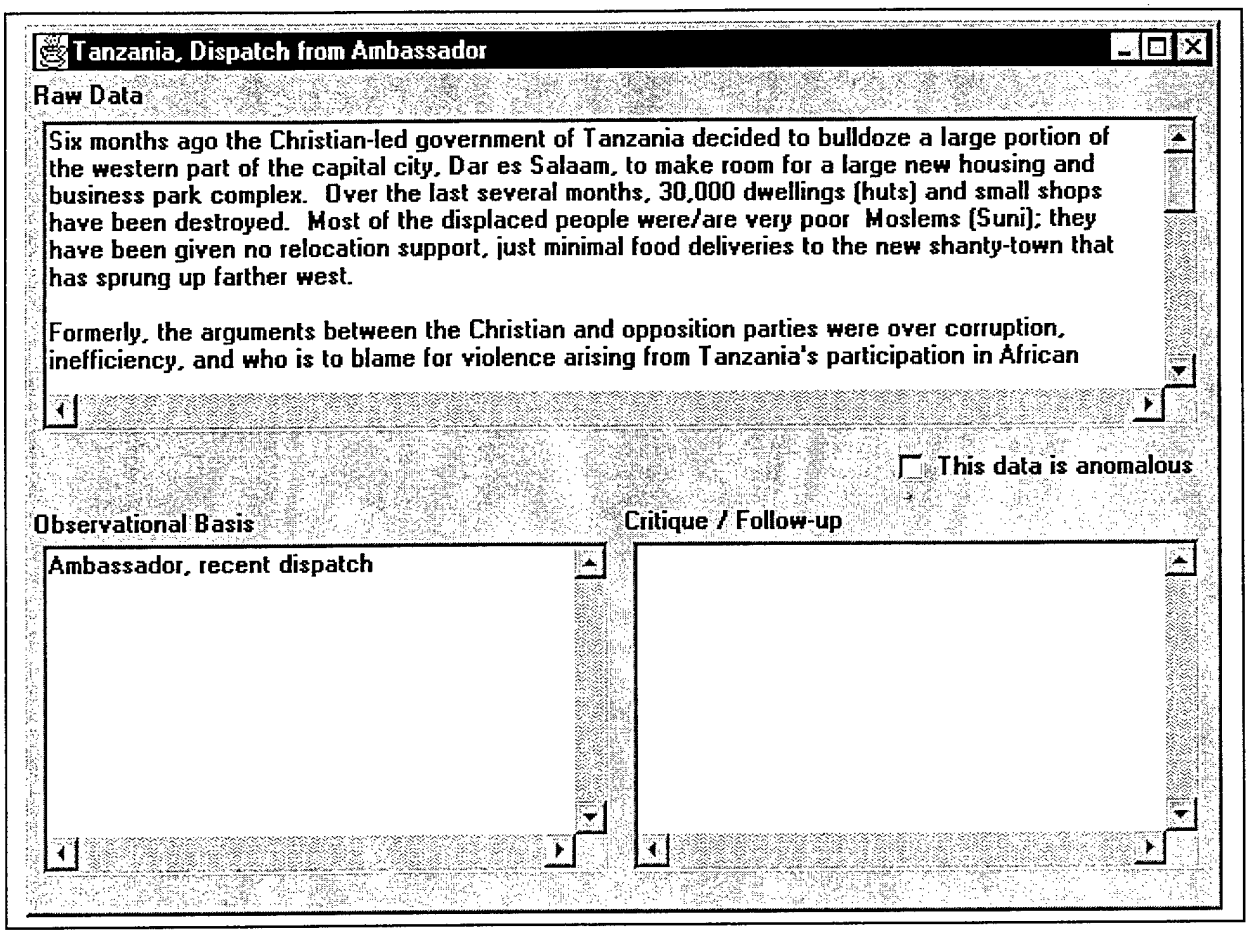

The following figure depicts the practices and social structure associated with the emerging organizational model. The one below it illustrates a refinement of the social structure to show the roles played by different agents.

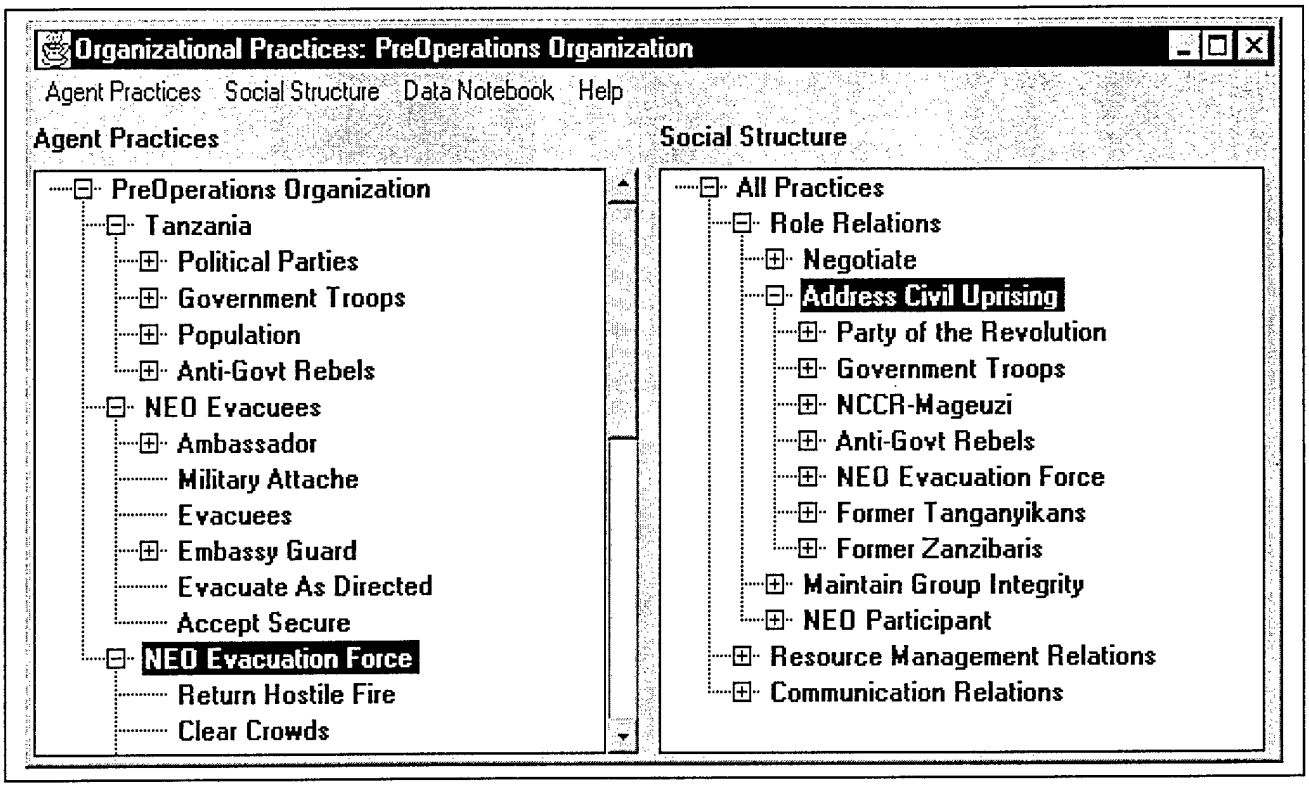

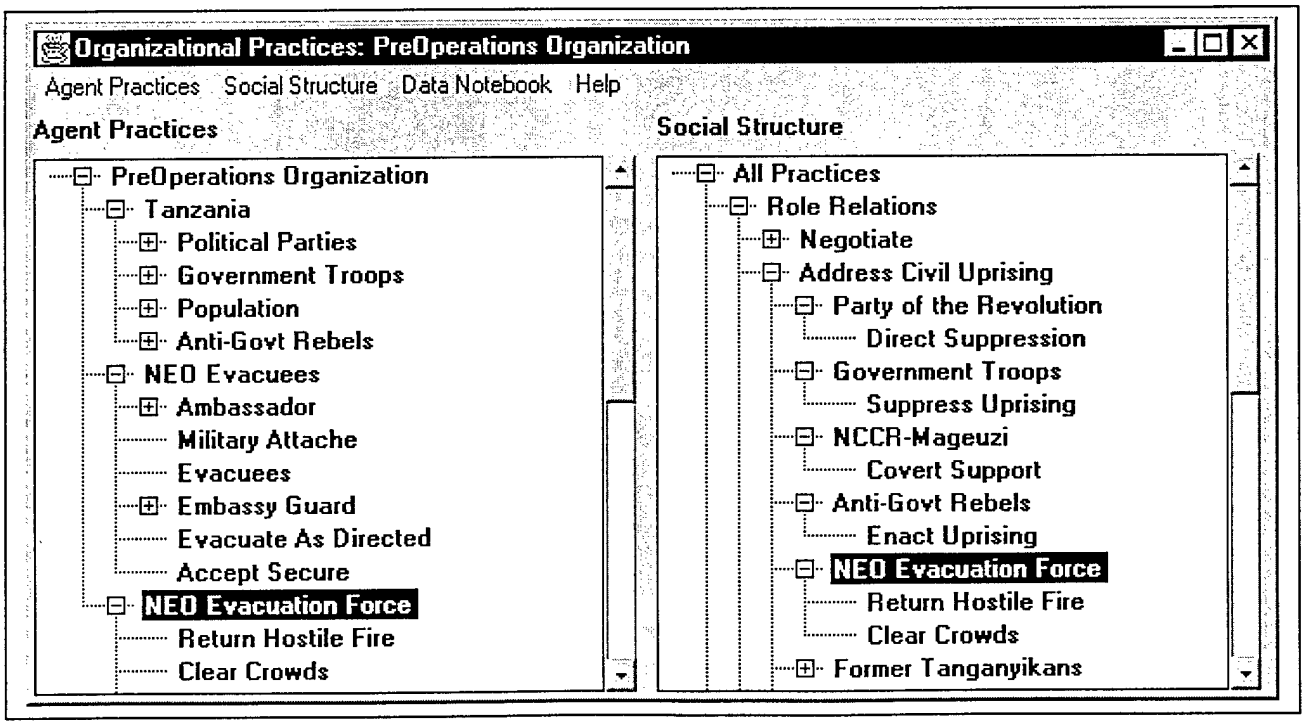

The following figure shows the way the anti-government rebels are modeled. The next figure will document the modeler's reasons for doing it this way.

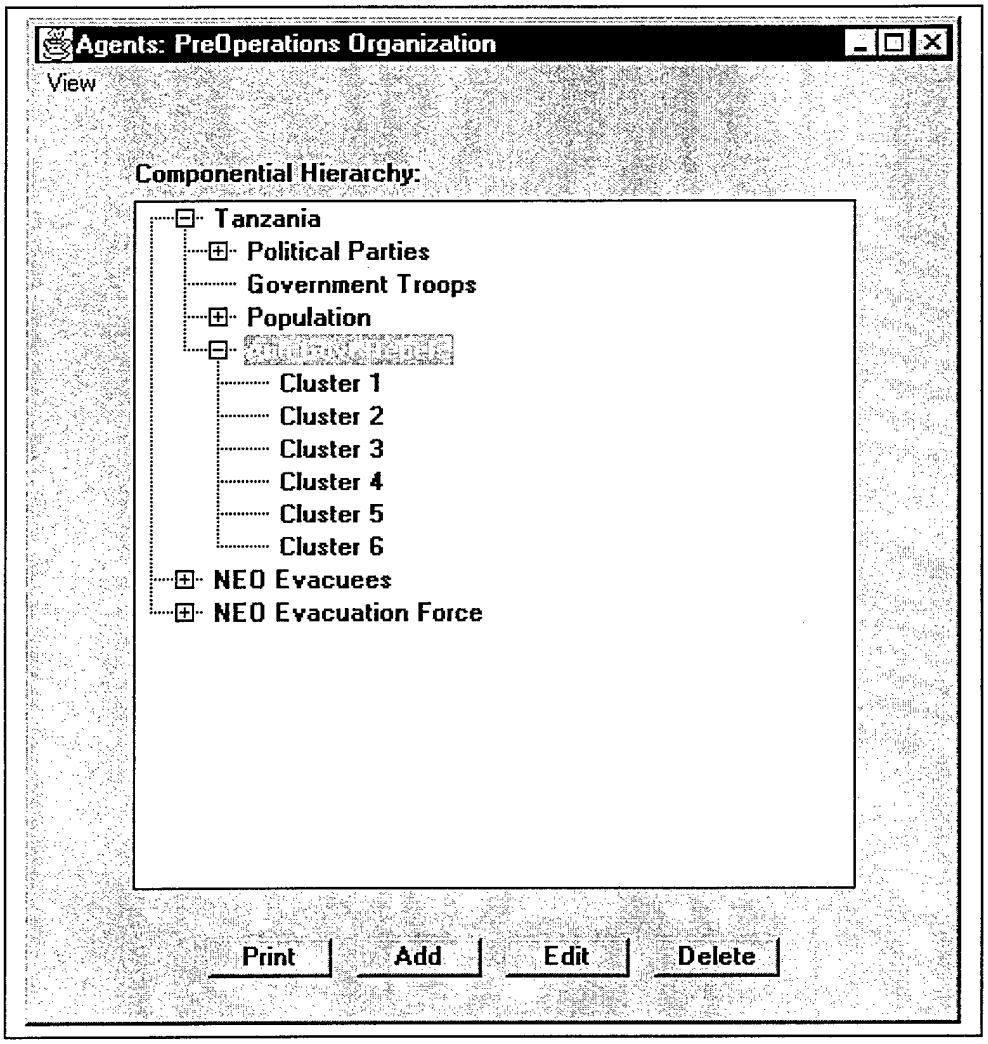

The following figure displays the rationale for the modeling decision. In particular, we illustrate the modeler being concerned about testing whether the evacuation can be completed prior to interference from the rebels (i.e., driving through their advance on the government offices risks incidents and casualties). Further, the modeler wants a simple way to note that they do not advance as a group, or as a military unit, but rather in a 'bloblike' manner. Hence, the invention of clusters that will have assigned statistical distributions for a motion model.

We also show the studies this model supports, and its link to data.

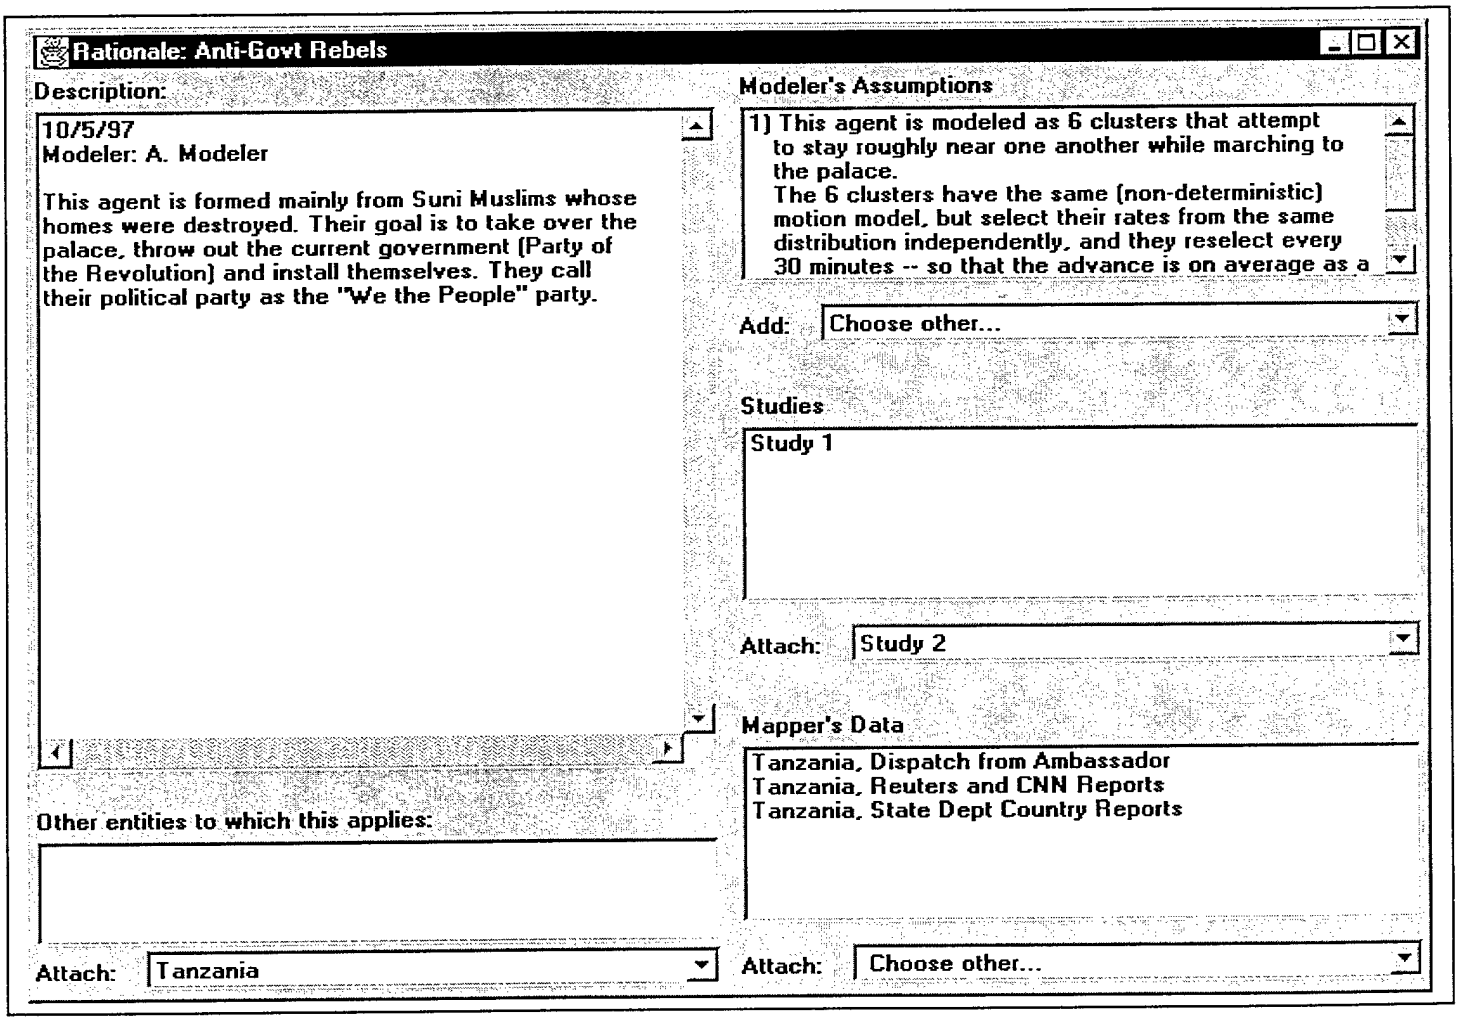

Here we show a practice that has been formalized as an action schema. It would be managed within the ambiguity reduction architecture. This is for the practice "SmallSetBattle" which manages what the clusters do when all of the rebels reach the goal (government offices). The initiator is reaching it; the binding environment is the data that is available; monitors test for failure, success, and suspend actions, and the procedure lays out the action plan.

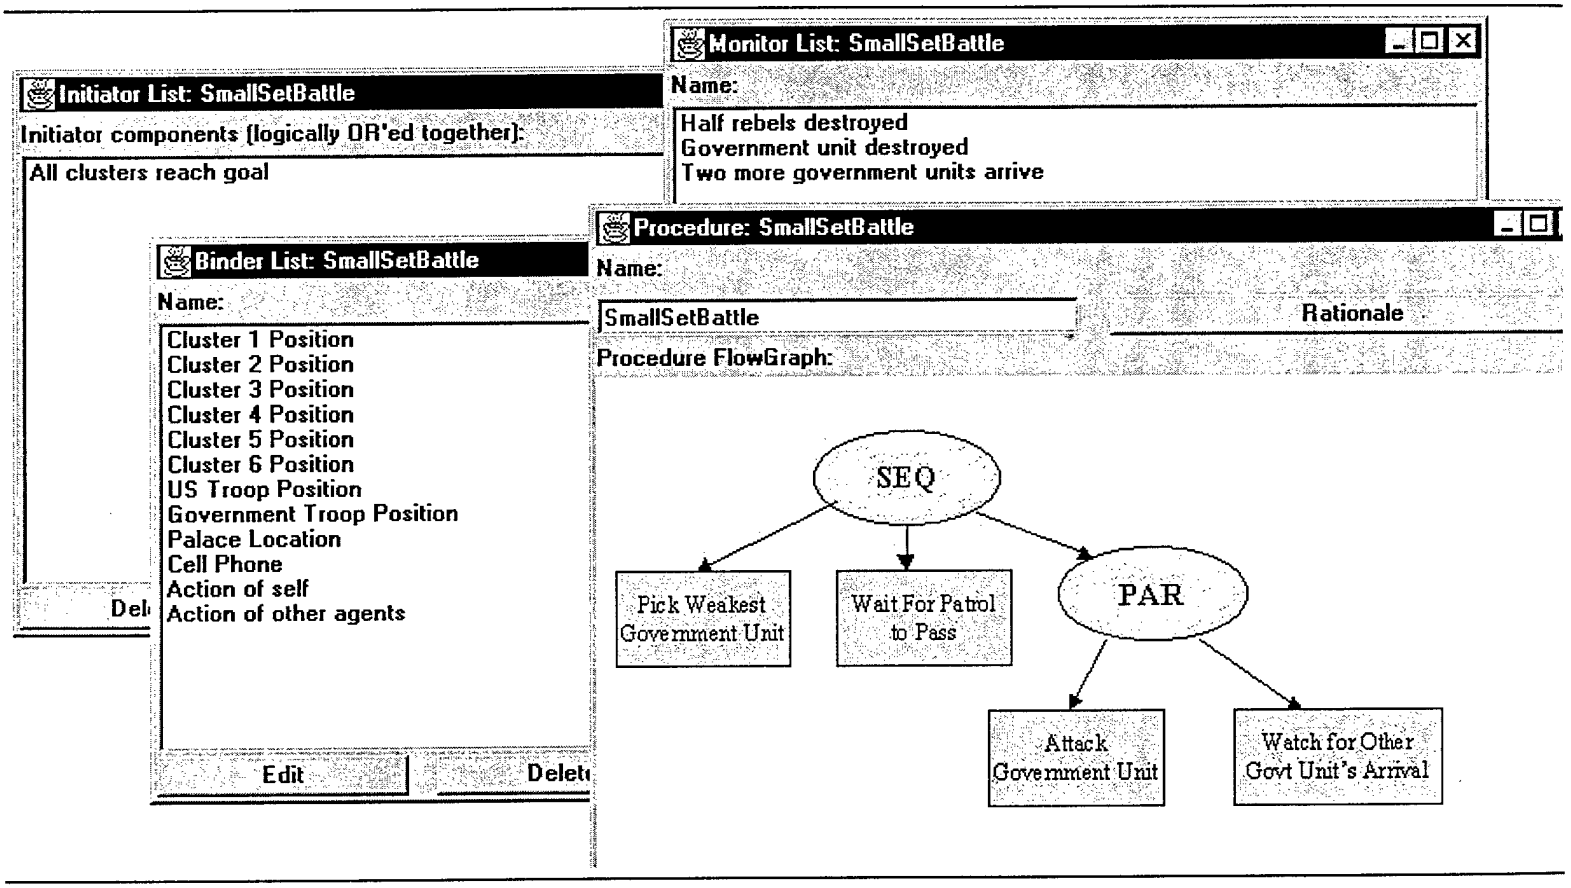

Here we show the main analysis interface, to support projective reasoning for validation and planning. On the left we show the kinds of analyses that are available (studies are projective tests for planning purposes, experiments are validation exercises) as well as the execution controls.

This picture also embodies the analysis discipline encoded into the technology. We depict an experiment as a topic, codified as a falsifiable hypothesis. Analysts name the structure than can vary (in this case, rate of advance of the rebels); a general protocol is provided to trace through a model and select structures that will be varied. This can range from replacing an entire agent, to some set of practices in an agent, to an element of a practice. Once variation structures are selected, analysts define the space over which variation occurs (in this case, a family probability distributions over rates of advance) – these are the cases or scenarios that are investigated via monte-carlo simulation. Finally, analysts introduce measurements to capture data during the run (e.g., when the evacuation troops cross paths with the rebels).

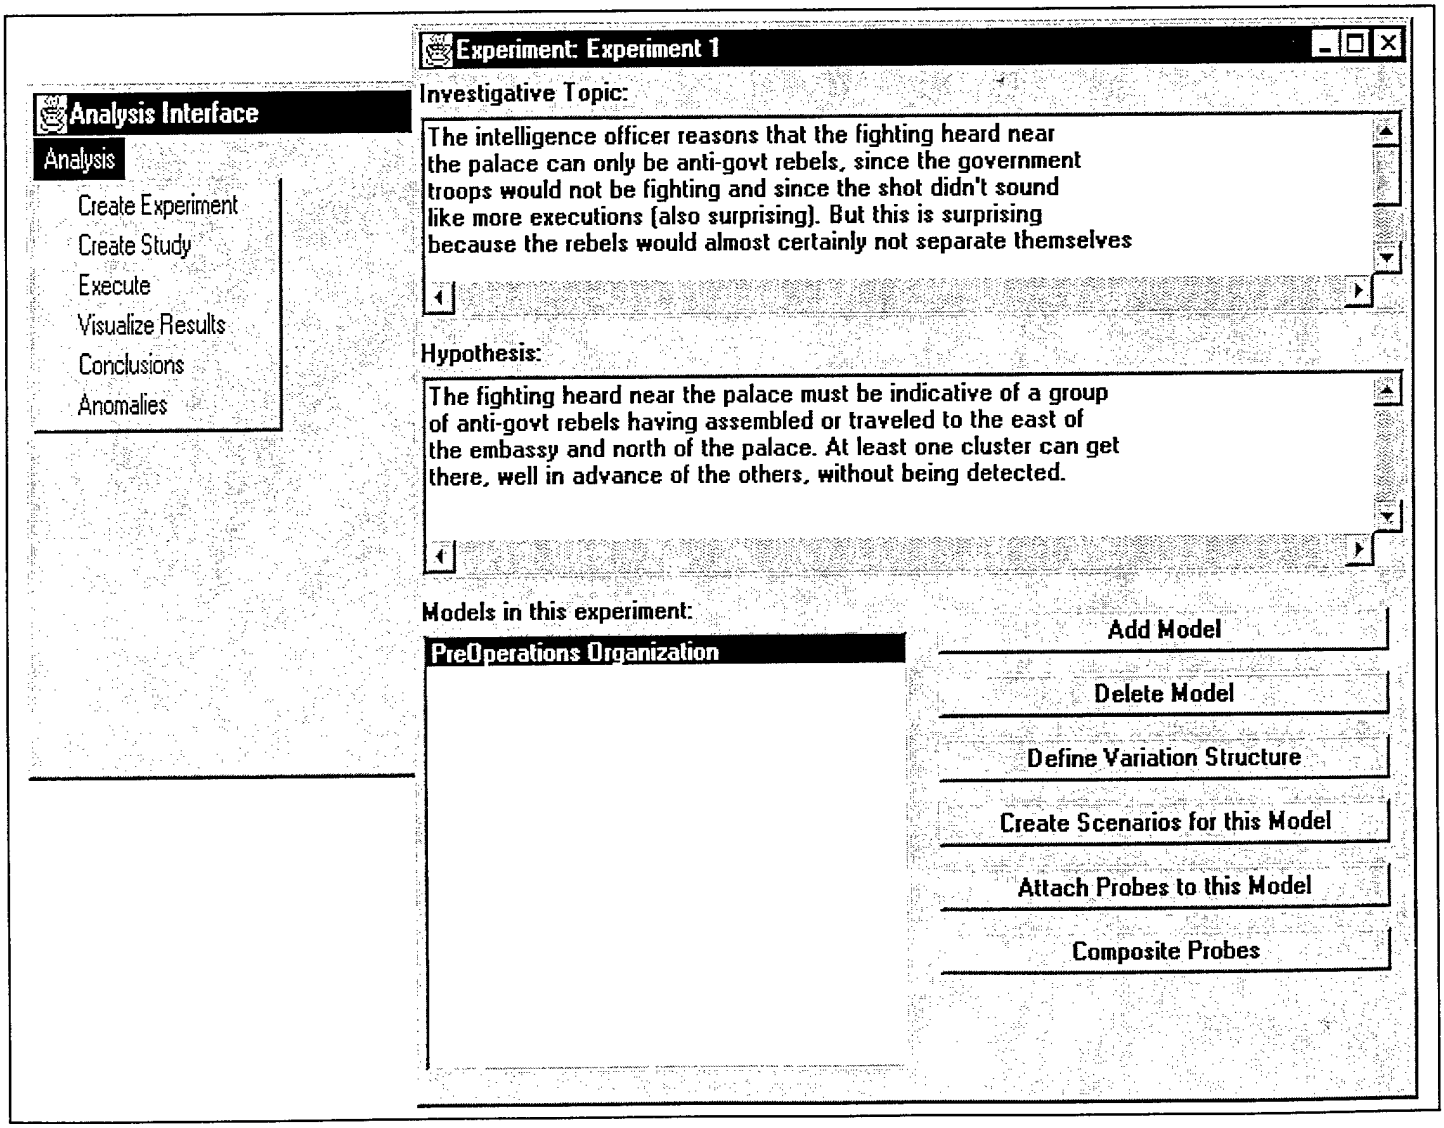

Here we show a picture of an anomaly. In this case, we have data that suggests there are agents near the government offices, and they are shooting. But in an experiment it is found that our model of the rebels is not compatible with them being there, and certainly not with some of them having gotten there (and so crossing the evacuation path) without evacuation troops seeing them. The experiment is similar to the previous one, but spreads out the distributions on motion as a way of testing whether one agent (that happens to sample fast rates of advance more frequently than others) gets in front of the other clusters. Interestingly, this time the measures are not only on the location of each cluster as a function of time, but also on the sensing/data of the evacuation troops. This is an illustration of the need for agent architectures – to capture not only how practices are managed but also the sensing and effecting capabilities of agents.

In short, this anomaly shows that an implication of the data contradicts an implication of our models. The representation of this contradiction, and the rationale for each of its two parts, formalizes the anomaly. Later, this and other anomalies lead to the breakdown of the model. At the same time, they contain data that suggests what new information to collect, and how to interpret it to create new models.

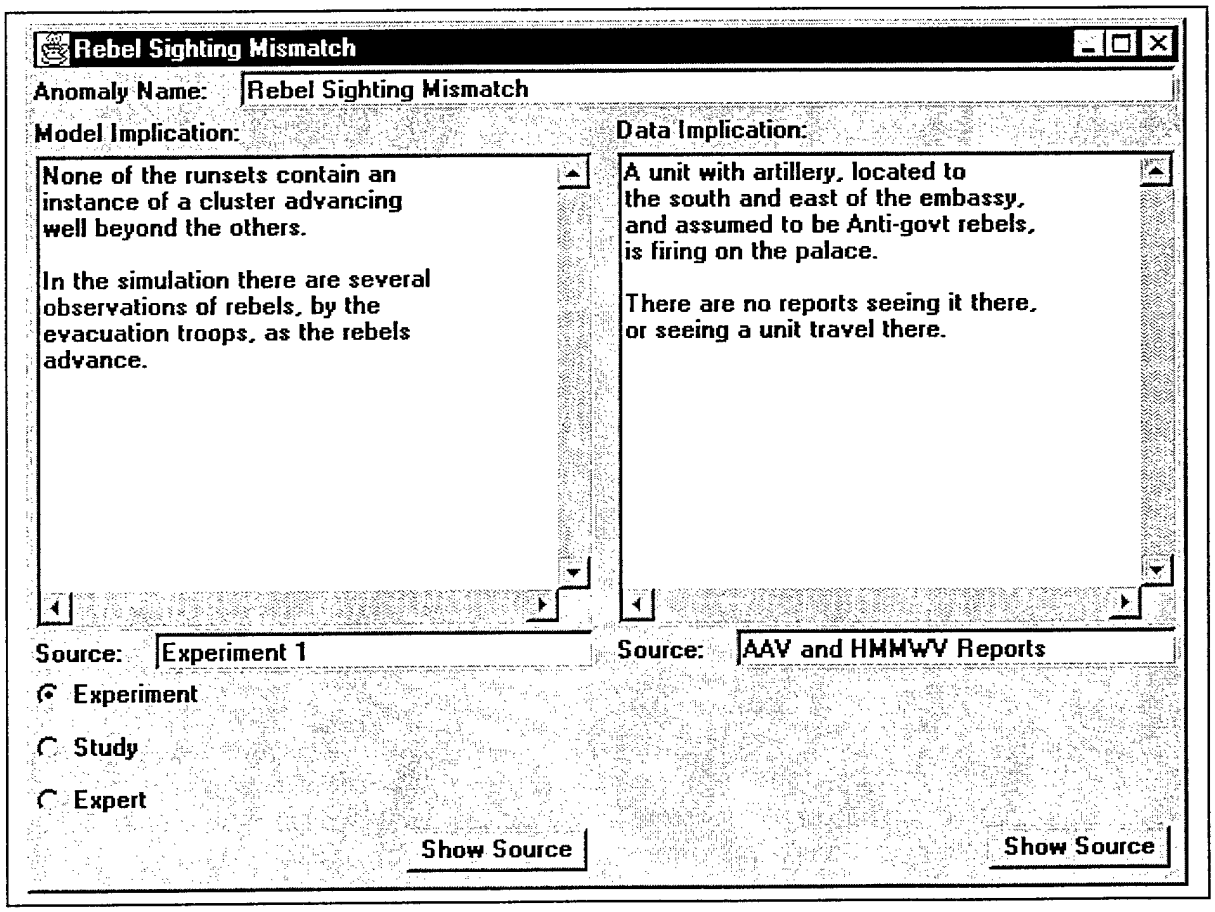

Schemes: the Underlying Computational Primitive for ACCORD and the CFA

We close our discussion of technology with a discussion of a unifying computational construct we have invented and that that underwrites our simulation, practice, agent, and CFA architecture models.

This construct, called a **scheme**, implements the special control regime required to reify our action schema construct, and so is used to instantiate models of practices. However, this construct is sufficiently general as to permit the reification of agents, organizations, and even the distributed simulation driver. Finally, we anticipate that it will be used to instantiate the CFA when we implement it.

The scheme construct is sufficient because each of these structures (practice, agent, organization, simulation driver, CFA) is in essence an action capacity whose computational evolution is made to depend on the occurrence (or non-occurrence) and specific content of information that arises during execution. In other words, a scheme is a feedback or closed-loop mechanism that provides for adaptation to interrupts of arbitrary computational grain-size. We have prototyped this construct and have implemented sample practices, agents, organizations, as well as a discrete simulation driver that employs multiple hardware processors.

A scheme is specified by the coupled-process specification. One process specification describes an action plan and an associated information access environment (patterns that the action plan might be influenced by). Action plans are themselves defined as a collection of continuations.[17] The executing action plan, and other executing practices or sensors, may alter the contents of the information environment.

The other process specification describes a collection of special structures, called monitors, designed to recognize whenever a pattern in the information environment warrants a change in the execution status of the action plan. This includes patterns that trigger it to start, suspend or abort or restart execution, or move the program counter from one continuation to another (i.e., 'advise' execution).

Once triggered and given execution resources, the coupled processes execute concurrently yet asynchronously. In this way interruptability of the executing action plan is achieved at arbitrary points.

---

[17] A program can be formalized in the usual manner as a having a starting configuration (configurations are descriptions of register states), a set of inputs that can be accepted, and etc. A computation is a sequence of instructions, and a trace is a set of configurations that arise as instructions are executed – instructions map configurations into subsequent ones. When a program is written as a set of continuations it allows, in addition to the starting configuration, a nonempty subset of the total set of configurations to be independently treated as starting configurations – each of which has its own range of acceptable inputs. This type of programming allows for arbitrary restarts, necessary to a general adaptive capacity.

Schemes can embody other schemes in their action plans. As such practices are implemented directly in schemes, and agents and organizations are recursively embedded systems of schemes.

Distribution of this Report

The following individuals and organizations comprise the distribution list for annual reports on this contract.

DARPA/ISO
Attn: Mr. Larry Willis
3701 N. Fairfax Dr.
Arlington, VA 22203-1714

DARPA/OASB
Attn: Library
3701 N. Fairfax Dr.
Arlington, VA 22203-1714

Defense Technical Information Center
Attn: OCC
8725 John J. Kingman Rd., Suite 0944
Ft. Belvoir, VA 22060-6218

DARPA/OASB
Attn: Ms. Connie Jacobs, SBIR PM
3701 N. Fairfax Dr.
Arlington, VA 22203-1714